Routing Options Feature Guide for QFabric Systems

Modified: 2017-01-30

Copyright © 2017, Juniper Networks, Inc.

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Routing Options Feature Guide for QFabric Systems
Copyright © 2017, Juniper Networks, Inc.
All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at http://www.juniper.net/support/eula.html. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About the Documentation xiii
Documentation and Release Notes xiii
Supported Platforms xiii
Using the Examples in This Manual xiii
Merging a Full Example xiv
Merging a Snippet xiv
Documentation Conventions xv
Documentation Feedback xvii
Requesting Technical Support xvii
Self-Help Online Tools and Resources xvii
Opening a Case with JTAC xviii

Part 1 Routing Options

Chapter 1 **Configuring Routing Options 3**

Overview of Routing Options 3
Configuring Static Routing 4
Configuring Per-Packet Load Balancing 4
Examples: Configuring Per-Packet Load Balancing 6
Understanding Virtual Router Routing Instances 7
Configuring Virtual Router Routing Instances 8
Understanding Distributed Periodic Packet Management 9
Configuring Distributed Periodic Packet Management 9
Disabling or Enabling Distributed Periodic Packet Management Globally . . . 10
Disabling or Enabling Distributed Periodic Packet Management for LACP Packets 10
Understanding Bidirectional Forwarding Detection (BFD) 11
Examples: Configuring BFD for Static Routes 11
Understanding BFD for Static Routes for Faster Network Failure Detection 11
Example: Configuring BFD for Static Routes for Faster Network Failure Detection 16
Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection 22
Example: Configuring BFD Authentication for Static Routes 28
Understanding BFD Authentication for Static Route Security 28
BFD Authentication Algorithms 29
Security Authentication Keychains 30
Strict Versus Loose Authentication 30
Example: Configuring BFD Authentication for Securing Static Routes 30

Copyright © 2017, Juniper Networks, Inc.

Understanding the Unified Forwarding Table . . . 36
Using the Unified Forwarding Table to Optimize Address Storage . . . 36
MAC Address and Host Address Memory Allocation . . . 36
LPM Table Memory Allocation . . . 38
Configuring the Unified Forwarding Table . . . 39
Configuring a Unified Forwarding Table Profile . . . 40
Configuring the Memory Allocation for Longest Prefix Match Entries . . . 41
Configuring the LPM Table With Junos OS Releases 13.2X51-D10 and 13.2X52-D10 . . . 41
Configuring the LPM Table With Junos OS Release 13.2x51-D15 and Later . . . 42
Monitoring Routing Information . . . 45
Verifying That Virtual Router Routing Instances Are Working . . . 47
Troubleshooting Virtual Routing Instances . . . 48
Direct Routes Not Leaked Between Routing Instances . . . 48

Part 2 Configuration Statements and Operational Commands

Chapter 2 Routing Options Configuration Statements . . . 53

active . . . 56
aggregate (Routing) . . . 57
as-path (Routing Options) . . . 59
autonomous-system . . . 61
backup-pe-group . . . 63
backups . . . 64
bandwidth . . . 65
bfd-liveness-detection (Routing Options Static Route) . . . 66
bgp-orf-cisco-mode . . . 70
bmp . . . 72
brief . . . 74
centralized . . . 75
community (Routing Options) . . . 76
confederation . . . 78
disable (Routing Options) . . . 79
description (Routing Instances) . . . 79
discard . . . 80
export (Routing Options) . . . 81
export-rib . . . 82
fate-sharing . . . 85
flow . . . 86
flow-map . . . 87
forwarding-cache (Flow Maps) . . . 88
forwarding-cache (Multicast) . . . 88
forwarding-options (chassis) . . . 89
forwarding-table . . . 92
generate . . . 93
graceful-restart (Enabling Globally) . . . 95
import . . . 96
import-policy . . . 97

Copyright © 2017, Juniper Networks, Inc.

import-rib . . . 98
indirect-next-hop . . . 99
install (Routing Options) . . . 100
instance-export . . . 101
instance-import . . . 101
instance-type . . . 102
interface (Multicast Static Routes) . . . 103
interface (Routing Instances) . . . 104
interface (Routing Options) . . . 105
interface-routes . . . 106
local-address (Routing Options) . . . 107
martians . . . 108
maximum-bandwidth . . . 109
maximum-paths . . . 110
maximum-prefixes . . . 112
med-igp-update-interval . . . 113
metric (Aggregate, Generated, or Static Route) . . . 114
multicast (Routing Options) . . . 115
no-qos-adjust . . . 116
num-65-127-prefix . . . 117
options (Routing Options) . . . 118
pim-to-igmp-proxy . . . 119
pim-to-mld-proxy . . . 120
policy (Aggregate and Generated Routes) . . . 121
policy (Flow Maps) . . . 122
policy-options . . . 123
policy-statement . . . 124
ppm . . . 128
ppm (Ethernet Switching) . . . 129
preference (Routing Options) . . . 130
prefix . . . 131
prefix-65-127-disable . . . 132
protocols . . . 133
qualified-next-hop (Static Routes) . . . 136
readvertise . . . 138
redundant-sources . . . 139
resolution . . . 140
resolution-ribs . . . 141
resolve . . . 142
restart-duration (Routing Options) . . . 143
retain . . . 144
reverse-oif-mapping . . . 145
rpf-check-policy (Routing Options RPF) . . . 146
rib (General) . . . 147
rib (Route Resolution) . . . 149
rib-group (Routing Options) . . . 150
rib-groups . . . 151
route-distinguisher-id . . . 153
route-record . . . 154

Copyright © 2017, Juniper Networks, Inc.

router-id . . . 155
routing-instances . . . 156
routing-options . . . 157
scope . . . 157
scope-policy . . . 158
source (Source-Specific Multicast) . . . 159
source-routing . . . 160
ssm-groups . . . 161
ssm-map (Routing Options Multicast) . . . 162
static (Routes) . . . 163
subscriber-leave-timer . . . 165
tag (Routing Options) . . . 166
threshold (Multicast Forwarding Cache) . . . 167
timeout (Flow Maps) . . . 168
timeout (Multicast) . . . 169
traceoptions . . . 170
upstream-interface . . . 173

Chapter 3 Routing Options Operational Commands . . . 175

clear ipv6 neighbors . . . 177
show as-path . . . 178
show as-path domain . . . 182
show as-path summary . . . 185
show ipv6 neighbors . . . 187
show ipv6 router-advertisement . . . 189
show route . . . 192
show route active-path . . . 200
show route all . . . 205
show route aspath-regex . . . 207
show route best . . . 209
show route brief . . . 212
show route community . . . 214
show route community-name . . . 216
show route damping . . . 218
show route detail . . . 223
show route exact . . . 243
show route export . . . 245
show route extensive . . . 248
show route flow validation . . . 266
show route forwarding-table . . . 268
show route inactive-path . . . 282
show route inactive-prefix . . . 286
show route instance . . . 288
show route label . . . 296
show route label-switched-path . . . 300
show route martians . . . 302
show route next-hop . . . 304
show route no-community . . . 310
show route protocol . . . 313

Copyright © 2017, Juniper Networks, Inc.

show route range 325
show route receive-protocol 329
show route resolution 338
show route snooping 341
show route source-gateway 344
show route summary 350
show route table 354
show route terse 383

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

List of Figures

Part 1 Routing Options

Chapter 1 Configuring Routing Options 3

Figure 1: Customer Routes Connected to a Service Provider 17
Figure 2: BFD Enabled on Qualified Next Hops 22
Figure 3: Customer Routes Connected to a Service Provider 31

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

List of Tables

About the Documentation xiii

Table 1: Notice Icons xv
Table 2: Text and Syntax Conventions xv

Part 1 Routing Options

Chapter 1 Configuring Routing Options 3

Table 3: Unified Forwarding Table Profiles on QFX3500, QFX3600, and QFX5100 Switches 37
Table 4: Unified Forwarding Table Profiles on QFX5200 Switches 37
Table 5: Example Host Table Combinations Using l2-profile-one on a QFX5100 Switch 38
Table 6: Example LPM Table Combinations Using L2 and L3 profiles With Junos OS 13.2X51-D10 and 13.2X52-D10 42
Table 7: LPM Table Combinations for L2 and L3 profiles With Junos OS 13.2X51-D15 and Later 43
Table 8: lpm-profile with unicast-in-lpm Option for QFX3500, QFX3600, and QFX5100 Switches 45
Table 9: Filtering Route Messages 46
Table 10: Summary of Key Routing Information Output Fields 46

Part 2 Configuration Statements and Operational Commands

Chapter 2 Routing Options Configuration Statements 53

Table 11: Unified Forwarding Table Profiles 91

Chapter 3 Routing Options Operational Commands 175

Table 12: show as-path Output Fields 179
Table 13: show as-path domain Output Fields 182
Table 14: show as-path summary Output Fields 185
Table 15: show ipv6 neighbors Output Fields 187
Table 16: show ipv6 router-advertisement Output Fields 189
Table 17: show route Output Fields 193
Table 18: show route damping Output Fields 219
Table 19: show route detail Output Fields 223
Table 20: Next-hop Types Output Field Values 229
Table 21: State Output Field Values 231
Table 22: Communities Output Field Values 233
Table 23: show route export Output Fields 245
Table 24: show route extensive Output Fields 248
Table 25: show route flow validation Output Fields 266

Copyright © 2017, Juniper Networks, Inc.

Table 26: show route forwarding-table Output Fields 271
Table 27: show route instance Output Fields 289
Table 28: show route martians Output Fields 302
Table 29: show route receive-protocol Output Fields 330
Table 30: show route resolution Output Fields 339
Table 31: show route summary Output Fields 350
Table 32: show route table Output Fields 355
Table 33: Next-hop Types Output Field Values 360
Table 34: State Output Field Values 362
Table 35: Communities Output Field Values 364
Table 36: show route terse Output Fields 383

Copyright © 2017, Juniper Networks, Inc.

About the Documentation

- Documentation and Release Notes on page xiii
- Supported Platforms on page xiii
- Using the Examples in This Manual on page xiii
- Documentation Conventions on page xv
- Documentation Feedback on page xvii
- Requesting Technical Support on page xvii

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at http://www.juniper.net/techpubs/.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at http://www.juniper.net/books.

Supported Platforms

For the features described in this document, the following platforms are supported:

- QFabric System

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

Copyright © 2017, Juniper Networks, Inc.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   system {
     scripts {
       commit {
         file ex-script.xsl;
       }
     }
   }
   interfaces {
     fxp0 {
       disable;
       unit 0 {
         family inet {
           address 10.0.0.1/24;
         }
       }
     }
   }
   ```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

   ```
   [edit]
   user@host# load merge /var/tmp/ex-script.conf
   load complete
   ```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

   For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

   ```
   commit {
     file ex-script-snippet.xsl; }
   ```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:

Copyright © 2017, Juniper Networks, Inc.

```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see CLI Explorer.

Documentation Conventions

Table 1 on page xv defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xv defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**

Copyright © 2017, Juniper Networks, Inc.

Table 2: Text and Syntax Conventions *(continued)*

Convention	Description	Examples
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric** ***metric*****>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [** ***community-ids*** **]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.

Copyright © 2017, Juniper Networks, Inc.

Table 2: Text and Syntax Conventions *(continued)*

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at http://www.juniper.net/techpubs/index.html, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at http://www.juniper.net/techpubs/feedback/.
- E-mail—Send your comments to **techpubs-comments@juniper.net**. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit http://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/

Copyright © 2017, Juniper Networks, Inc.

- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: http://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://tools.juniper.net/SerialNumberEntitlementSearch/

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at http://www.juniper.net/cm/.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see http://www.juniper.net/support/requesting-support.html.

Copyright © 2017, Juniper Networks, Inc.

PART 1

Routing Options

- Configuring Routing Options on page 3

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 1

Configuring Routing Options

- Overview of Routing Options on page 3
- Configuring Static Routing on page 4
- Configuring Per-Packet Load Balancing on page 4
- Examples: Configuring Per-Packet Load Balancing on page 6
- Understanding Virtual Router Routing Instances on page 7
- Configuring Virtual Router Routing Instances on page 8
- Understanding Distributed Periodic Packet Management on page 9
- Configuring Distributed Periodic Packet Management on page 9
- Understanding Bidirectional Forwarding Detection (BFD) on page 11
- Examples: Configuring BFD for Static Routes on page 11
- Example: Configuring BFD Authentication for Static Routes on page 28
- Understanding the Unified Forwarding Table on page 36
- Configuring the Unified Forwarding Table on page 39
- Monitoring Routing Information on page 45
- Verifying That Virtual Router Routing Instances Are Working on page 47
- Troubleshooting Virtual Routing Instances on page 48

Overview of Routing Options

In addition to dynamic routing protocols, you can configure static routing on QFX Series switches. You can also configure a variety of protocol-independent routing properties, such as

- Per-packet load balancing (equal cost multipath routing)
- Autonomous system numbers
- Autonomous system confederation members
- Router identifiers
- Routing table groups
- Multicast scoping

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- Understanding Distributed Periodic Packet Management on page 9
- Understanding Virtual Router Routing Instances on page 7

Configuring Static Routing

Static routes are routes that are manually configured and entered into the routing table.

The switch uses static routes:

- When the switch does not have a route to a destination that has a better (lower) *preference* value. The preference is an arbitrary value in the range from 0 through 255 that the software uses to rank routes received from different protocols, interfaces, or remote systems. The routing protocol process generally determines the active route by selecting the route with the lowest preference value. In the given range, **0** is the lowest and **255** is the highest.
- When the switch cannot determine the route to a destination.
- When the switch is forwarding unroutable packets.

To configure basic static route options using the CLI:

- To configure the switch's default gateway:

  ```
  [edit]
  user@switch# set routing–options static route 0.0.0.0/0 next-hop 10.0.1.1
  ```

- To configure a static route and specify the next address to be used when routing traffic to the static route:

  ```
  [edit]
  user@switch# set routing-options static route 198.51.100.0/24 next-hop 10.0.0.2
  ```

- To always keep the static route in the forwarding table:

  ```
  [edit]
  user@switch# set routing-options static route 198.51.100.0/24 retain
  ```

- To prevent the static route from being readvertised:

  ```
  [edit]
  user@switch# set routing-options static route 198.51.100.0/24 no-readvertise
  ```

- To remove inactive routes from the forwarding table:

  ```
  [edit]
  user@switch# set routing-options static route 198.51.100.0/24 active
  ```

Related Documentation
- Monitoring Routing Information on page 45

Configuring Per-Packet Load Balancing

By default, when there are multiple equal-cost paths to the same destination for the active route, Junos OS chooses one of the next-hop addresses to install into the forwarding

Copyright © 2017, Juniper Networks, Inc.

table in a random fashion. Whenever the set of next hops for a destination changes in any way, the next-hop address is chosen again, also in a random fashion.

You can configure Junos OS so that, for the active route, all next-hop addresses for a destination are installed in the forwarding table. This is called per-packet load balancing. You can use this feature to spread traffic across multiple paths.

On a QFX3500 standalone switch, with static routing configured, whenever a route pointing to an ECMP next-hop changes to a new ECMP next-hop with a different member list but contains the exact member count as before, the location of the retained members in the new member list is the same as in the old member list.

For example, if you have the following configuration on the switch:

```
set routing-options static route 0.0.0.0/0 next-hop 192.0.2.2
set routing-options static route 0.0.0.0/0 next-hop 192.0.2.3
set routing-options static route 0.0.0.0/0 next-hop 192.0.2.4
set routing-options static route 0.0.0.0/0 next-hop 198.51.100.2
set routing-options static route 0.0.0.0/0 next-hop 198.51.100.3
set routing-options static route 0.0.0.0/0 next-hop 198.51.100.4
set routing-options static route 0.0.0.0/0 next-hop 203.0.113.1
set routing-options static route 0.0.0.0/0 next-hop 203.0.113.2
```

and want to change the first and eighth route to point to another location, you can issue the following commands:

```
delete routing-options static route 0.0.0.0/0 next-hop 192.0.2.3
delete routing-options static route 0.0.0.0/0 next-hop 203.0.113.2
set routing-options static route 0.0.0.0/0 next-hop 192.0.2.5
set routing-options static route 0.0.0.0/0 next-hop 203.0.113.3
```

This configuration does not affect the second next-hop through the seventh next-hop.

When per-packet load balancing is configured, traffic is divided into individual flows (up to a maximum of 16). Packets for an individual flow are sent out a single interface. To determine flows, the switch examines each of the following packet fields:

- Source IP address
- Destination IP address
- Protocol
- Source port number
- Destination port number
- Source interface index
- Type of service (ToS)

The switch recognizes packets in which all of these parameters are identical and ensures that these packets are sent out through the same interface. This prevents problems that might otherwise occur with packets arriving at their destination out of their original sequence.

Copyright © 2017, Juniper Networks, Inc.

NOTE: Load balancing is not supported on management interfaces.

The following steps shows how to configure per-packet load balancing:

1. Define a load-balancing routing policy by including one or more **policy-statement** statements at the **[edit policy-options]** hierarchy level, defining an action of **load-balance per-packet**:

```
policy-statement policy-name {
  from {
    match-conditions;
    route-filter destination-prefix match-type <actions>;
    prefix-list name;
  }
  then {
    load-balance per-packet;
  }
}
```

2. Apply the policy to routes exported from the routing table to the forwarding table. To do this, include the **forwarding-table** and **export** statements:

```
forwarding-table {
  export policy-name;
}
```

When you enable per-packet load balancing on a QFabric system, packets might be switched across the fabric even though there is a local port with a same-cost route to the destination. For example, if per-packet load balancing is enabled and a packet arrives at network Node device A, it might be switched to network Node device B and forwarded from there even if there is a same-cost route through a port on Node device A to the destination. In this case, traffic transits the fabric needlessly. you can configure a QFabric system to choose a locally switched route if one is available. To enable this feature, include the **ecmp-do-local-lookup** statement at the **[edit forwarding-options]** hierarchy level.

Related Documentation

- Examples: Configuring Per-Packet Load Balancing on page 6
- *Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups*

Examples: Configuring Per-Packet Load Balancing

Perform per-packet load balancing for all routes:

```
[edit]
policy-options {
  policy-statement load-balancing-policy {
    then {
      load-balance per-packet;
    }
  }
```


Copyright © 2017, Juniper Networks, Inc.

```
}
routing-options {
  forwarding-table {
    export load-balancing-policy;
  }
}
```

Perform per-packet load balancing for a limited set of routes:

```
[edit]
policy-options {
  policy-statement load-balancing-policy {
    from {
      route-filter 192.168.10/24 orlonger;
      route-filter 9.114/16 orlonger;
    }
    then {
      load-balance per-packet;
    }
  }
}
routing-options {
  forwarding-table {
    export load-balancing-policy;
  }
}
```

Related Documentation

- Configuring Per-Packet Load Balancing on page 4
- *Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups*

Understanding Virtual Router Routing Instances

Virtual router routing instances allow administrators to divide a QFX Series switch into multiple independent virtual routers, each with its own routing table. Virtual router routing instances enable you to isolate traffic without using multiple devices to segment your network. You can create routing instances for unicast routing protocols and PIM sparse mode.

Each virtual router routing instance consists of sets of the following:

- Routing tables
- Interfaces that belong to these routing tables
- Routing protocol configurations
- Routing option configurations

You can use virtual router routing instances to isolate customer traffic on a network and to bind customer-specific routing instances to customer-owned interfaces. Each interface can belong only to one routing instance.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Configuring Virtual Router Routing Instances on page 8

Configuring Virtual Router Routing Instances

Use virtual router routing instances to divide a QFX Series switch into multiple independent virtual routers, each with its own routing table. Virtual router routing instances enable you to isolate traffic without using multiple devices to segment your network. You can create routing instances for unicast routing protocols and PIM sparse mode.

To configure virtual router routing instances:

1. Create a routing instance:

   ```
   [edit routing-instances]
   user@switch# set routing-instance-name instance-type virtual-router
   ```

   NOTE: The default routing instance, **master**, refers to the main **inet.0** routing table. The master routing instance is reserved and cannot be specified as a routing instance.

2. Bind each routing instance to the corresponding interfaces:

   ```
   [edit routing-instances]
   user@switch# set routing-instance-name interface
   device-name:type-fpc/pic/port.logical-unit-number
   ```

   NOTE:
   - You must bind routing instances to interfaces from the Node devices assigned to the network Node group only. If you try to bind routing instances to interfaces from the Node devices assigned to server Node groups, the configuration does not commit.
   - You can bind an interface to one routing instance only.

3. Create each of the logical interfaces bound to each routing instance:

   ```
   [edit interfaces]
   user@switch# set device-name:type-fpc/pic/port unit logical-unit-number family inet address
   ip-address
   ```

   NOTE: Do not create a logical interface using the **family ethernet-switching** option in this step. Binding an interface using the **family ethernet-switching** option to a routing instance can cause the interface to shut down.

4. (Optional) Configure routing protocols for the routing instance at the **[edit routing-instances *routing-instance-name* protocols]** hierarchy level.
5. (Optional) Configure routing options for the routing instance at the **[edit routing-instances *routing-instance-name* routing-options]** hierarchy level.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- Understanding Virtual Router Routing Instances on page 7
- Verifying That Virtual Router Routing Instances Are Working on page 47

Understanding Distributed Periodic Packet Management

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks for particular processes so that other processes can more optimally direct their resources. PPM is responsible for the periodic transmission of packets on behalf of its various client processes, which include the processes that control the Link Aggregation Control Protocol (LACP) and Bidirectional Forwarding Detection (BFD) protocol, and also for receiving packets on behalf of these client processes. PPM also gathers some statistics and sends process-specific packets. PPM cannot be disabled and is always running on any operational switch.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and the access interfaces for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

You can disable distributed PPM for all protocols that use PPM. You can also disable distributed PPM for LACP packets only.

BEST PRACTICE: **We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.**

Related Documentation
- Configuring Distributed Periodic Packet Management on page 9

Configuring Distributed Periodic Packet Management

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks so that other processes can more optimally direct their resources.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and the access interfaces for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

You can disable distributed PPM for all protocols that use PPM. You can also disable distributed PPM for Link Aggregation Control Protocol (LACP) packets only.

Copyright © 2017, Juniper Networks, Inc.

BEST PRACTICE: We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

This topic describes:

- Disabling or Enabling Distributed Periodic Packet Management Globally on page 10
- Disabling or Enabling Distributed Periodic Packet Management for LACP Packets on page 10

Disabling or Enabling Distributed Periodic Packet Management Globally

Distributed PPM is enabled by default. Disable distributed PPM if you need to move all PPM processing to the Routing Engine. Enable distributed PPM if it was previously disabled and you need to run distributed PPM.

To disable distributed PPM:

```
[edit routing-options]
user@switch# set ppm no-delegate-processing
```

To enable distributed PPM if it was previously disabled:

```
[edit routing-options]
user@switch# delete ppm no-delegate-processing
```

Disabling or Enabling Distributed Periodic Packet Management for LACP Packets

Distributed PPM is enabled by default. Disable distributed PPM for only LACP packets if you need to move all PPM processing for LACP packets to the Routing Engine.

To disable distributed PPM for LACP packets:

```
[edit protocols]
user@switch# set lacp ppm centralized
```

To enable distributed PPM for LACP packets if it was previously disabled:

```
[edit protocols]
user@switch# delete lacp ppm centralized
```

Related Documentation

- Understanding Distributed Periodic Packet Management on page 9
- *Understanding Aggregated Ethernet Interfaces and LACP*

Copyright © 2017, Juniper Networks, Inc.

Understanding Bidirectional Forwarding Detection (BFD)

The Bidirectional Forwarding Detection (BFD) protocol is a simple mechanism that detects failures in a network and works in a wide variety of network environments and topologies. In BFD operation, switches exchanges BFD hello packets at a specified interval and detect a neighbor failure if they do not receive a reply after a specified interval. The BFD failure detection timers support shorter time limits than the static route failure detection mechanisms, so they can provide faster detection of failures.

To configure faster failure detection, use lower BFD timer values. The timers can automatically adapt to a higher value if an adjacency fails, and they also adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. In this case, a back-off algorithm increases the receive interval by two if the local BFD instance is the reason for the session flap and increases the transmission interval by two if the remote BFD instance is the reason for the session flap.

You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. This command is hitless, meaning that it does not affect traffic flow.

NOTE: QFX10000 series switches support BFD timer values of less than 1 second.

NOTE: BFD is not supported with ISIS for IPV6 on QFX10000 series switches.

NOTE: EX3300 supports BFD over static routes only.

Related Documentation

- Examples: Configuring BFD for Static Routes on page 11
- Example: Configuring BFD Authentication for Static Routes on page 28

Examples: Configuring BFD for Static Routes

- Understanding BFD for Static Routes for Faster Network Failure Detection on page 11
- Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 16
- Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection on page 22

Understanding BFD for Static Routes for Faster Network Failure Detection

The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures in a network. BFD works with a wide variety of network environments and topologies. A pair of routing devices exchanges BFD packets. Hello packets are sent

Copyright © 2017, Juniper Networks, Inc.

at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. The BFD failure detection timers have shorter time limits than the static route failure detection mechanisms, so they provide faster detection.

The BFD failure detection timers are adaptive and can be adjusted to be faster or slower. The lower the BFD failure detection timer value, the faster the failure detection and vice versa. For example, the timers can adapt to a higher value if the adjacency fails (that is, the timer detects failures more slowly). Or a neighbor can negotiate a higher value for a timer than the configured value. The timers adapt to a higher value when a BFD session flap occurs more than three times in a span of 15 seconds. A back-off algorithm increases the receive (Rx) interval by two if the local BFD instance is the reason for the session flap. The transmission (Tx) interval is increased by two if the remote BFD instance is the reason for the session flap. You can use the **clear bfd adaptation** command to return BFD interval timers to their configured values. The **clear bfd adaptation** command is hitless, meaning that the command does not affect traffic flow on the routing device.

By default, BFD is supported on single-hop static routes.

To enable failure detection, include the bfd-liveness-detection statement in the static route configuration.

NOTE: Starting with Junos OS Release 15.1X49-D70, the **bfd-liveness-detection** command includes the description field. The description is an attribute under the **bfd-liveness-detection** object and it is supported only on SRX Series devices. This field is applicable only for the static routes.

In Junos OS Release 9.1 and later, the BFD protocol is supported for IPv6 static routes. Global unicast and link-local IPv6 addresses are supported for static routes. The BFD protocol is not supported on multicast or anycast IPv6 addresses. For IPv6, the BFD protocol supports only static routes and only in Junos OS Release 9.3 and later. IPv6 for BFD is also supported for the eBGP protocol.

NOTE:

Inline BFD is supported on PTX5000 routers with third-generation FPCs starting in Junos OS Release 15.1F3 and 16.1R2. Inline BFD is supported on PTX3000 routers with third-generation FPCs starting in Junos OS Release 15.1F6 and 16.1R2.

There are three types of BFD sessions based on the source from which BFD packets are sent to the neighbors. Different types of BFD sessions and their descriptions are given in the table below:

Type of BFD session	Description
Non-distributed BFD	BFD sessions running completely on the Routing Engine.

Copyright © 2017, Juniper Networks, Inc.

Type of BFD session	Description
Distributed BFD	BFD sessions running on the Packet Forwarding Engine.
Inline BFD NOTE: Starting in Junos OS Release 13.3, inline BFD is supported only on static MX Series routers with MPCs/MICs that have configured **enhanced-ip**. NOTE: Starting in Junos OS Release 16.1R1, the inline BFD sessions are supported on integrated routing and bridging (IRB) interfaces.	BFD sessions running on the FPC hardware.

To configure the BFD protocol for IPv6 static routes, include the **bfd-liveness-detection** statement at the **[edit routing-options rib inet6.0 static route *destination-prefix*]** hierarchy level.

In Junos OS Release 8.5 and later, you can configure a hold-down interval to specify how long the BFD session must remain up before a state change notification is sent.

To specify the hold-down interval, include the **holddown-interval** statement in the BFD configuration.

You can configure a number in the range from 0 through 255,000 milliseconds. The default is 0. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

NOTE: If a single BFD session includes multiple static routes, the hold-down interval with the highest value is used.

To specify the minimum transmit and receive intervals for failure detection, include the **minimum-interval** statement in the BFD configuration.

This value represents both the minimum interval after which the local routing device transmits hello packets and the minimum interval after which the routing device expects to receive a reply from the neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

Copyright © 2017, Juniper Networks, Inc.

NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with NSR configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

NOTE: SRX Series devices do not support distributed BFD.

To specify the minimum receive interval for failure detection, include the **minimum-receive-interval** statement in the BFD configuration. This value represents the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a number in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route *destination-prefix* bfd-liveness-detection]** hierarchy level.

To specify the number of hello packets not received by the neighbor that causes the originating interface to be declared down, include the **multiplier** statement in the BFD configuration.

The default value is 3. You can configure a number in the range from 1 through 255.

To specify a threshold for detecting the adaptation of the detection time, include the **threshold** statement in the BFD configuration.

When the BFD session detection time adapts to a value equal to or higher than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the **minimum-interval** or the **minimum-receive-interval** value. The threshold must be a higher value than the multiplier for either of these configured values. For example if the **minimum-receive-interval** is 300 ms and the **multiplier** is 3, the total detection time is 900 ms. Therefore, the detection time threshold must have a value higher than 900.

Copyright © 2017, Juniper Networks, Inc.

To specify the minimum transmit interval for failure detection, include the **transmit-interval minimum-interval** statement in the BFD configuration.

This value represents the minimum interval after which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route *destination-prefix* bfd-liveness-detection]** hierarchy level.

To specify the threshold for the adaptation of the transmit interval, include the **transmit-interval threshold** statement in the BFD configuration.

The threshold value must be greater than the transmit interval. When the BFD session transmit time adapts to a value greater than the threshold, a single trap and a system log message are sent. The detection time is based on the multiplier of the value for the **minimum-interval** or the **minimum-receive-interval** statement at the **[edit routing-options static route *destination-prefix* bfd-liveness-detection]** hierarchy level. The threshold must be a higher value than the multiplier for either of these configured values.

To specify the BFD version, include the **version** statement in the BFD configuration. The default is to have the version detected automatically.

To include an IP address for the next hop of the BFD session, include the **neighbor** statement in the BFD configuration.

NOTE: You must configure the neighbor statement if the next hop specified is an interface name. If you specify an IP address as the next hop, that address is used as the neighbor address for the BFD session.

In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions.

To disable BFD adaptation, include the **no-adaptation** statement in the BFD configuration.

NOTE: We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation in your network.

NOTE: If BFD is configured only on one end of a static route, the route is removed from the routing table. BFD establishes a session when BFD is configured on both ends of the static route.

BFD is not supported on ISO address families in static routes. BFD does support IS-IS.

If you configure graceful Routing Engine switchover (GRES) at the same time as BFD, GRES does not preserve the BFD state information during a failover.

Copyright © 2017, Juniper Networks, Inc.

Example: Configuring BFD for Static Routes for Faster Network Failure Detection

This example shows how to configure Bidirectional Forwarding Detection (BFD) for static routes.

- Requirements on page 16
- Overview on page 16
- Configuration on page 17
- Verification on page 20

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

There are many practical applications for static routes. Static routing is often used at the network edge to support attachment to stub networks, which, given their single point of entry and egress, are well suited to the simplicity of a static route. In Junos OS, static routes have a global preference of 5. Static routes are activated if the specified next hop is reachable.

In this example, you configure the static route 192.168.47.0/24 from the provider network to the customer network, using the next-hop address of 172.16.1.2. You also configure a static default route of 0.0.0.0/0 from the customer network to the provider network, using a next-hop address of 172.16.1.1.

For demonstration purposes, some loopback interfaces are configured on Device B and Device D. These loopback interfaces provide addresses to ping and thus verify that the static routes are working.

Figure 1 on page 17 shows the sample network.

Copyright © 2017, Juniper Networks, Inc.

Figure 1: Customer Routes Connected to a Service Provider

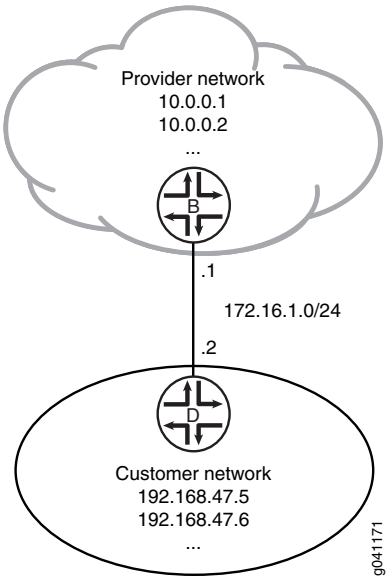

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device B

```
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
    1000
set routing-options static route 192.168.47.0/24 bfd-liveness-detection description
    Site-xxx
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
```

Device D

```
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
```

Copyright © 2017, Juniper Networks, Inc.

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD for static routes:

1. On Device B, configure the interfaces.

   ```
   [edit interfaces]
   user@B# set ge-1/2/0 unit 0 description B->D
   user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24
   user@B# set lo0 unit 57 family inet address 10.0.0.1/32
   user@B# set lo0 unit 57 family inet address 10.0.0.2/32
   ```

2. On Device B, create a static route and set the next-hop address.

   ```
   [edit routing-options]
   user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2
   ```

3. On Device B, configure BFD for the static route.

   ```
   [edit routing-options]
   user@B# set static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
     1000
   set routing-options static route 192.168.47.0/24 bfd-liveness-detection description
     Site-xxx
   ```

4. On Device B, configure tracing operations for BFD.

   ```
   [edit protocols]
   user@B# set bfd traceoptions file bfd-trace
   user@B# set bfd traceoptions flag all
   ```

5. If you are done configuring Device B, commit the configuration.

   ```
   [edit]
   user@B# commit
   ```

6. On Device D, configure the interfaces.

   ```
   [edit interfaces]
   user@D# set ge-1/2/0 unit 1 description D->B
   user@D# set ge-1/2/0 unit 1 family inet address 172.16.1.2/24
   user@D# set lo0 unit 2 family inet address 192.168.47.5/32
   user@D# set lo0 unit 2 family inet address 192.168.47.6/32
   ```

7. On Device D, create a static route and set the next-hop address.

   ```
   [edit routing-options]
   user@D# set static route 0.0.0.0/0 next-hop 172.16.1.1
   ```

8. On Device D, configure BFD for the static route.

   ```
   [edit routing-options]
   user@D# set static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
   ```

9. On Device D, configure tracing operations for BFD.

   ```
   [edit protocols]
   user@D# set bfd traceoptions file bfd-trace
   user@D# set bfd traceoptions flag all
   ```


Copyright © 2017, Juniper Networks, Inc.

10. If you are done configuring Device D, commit the configuration.

```
[edit]
user@D# commit
```

Results

Confirm your configuration by issuing the **show interfaces**, **show protocols**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Device B

```
user@B# show interfaces
ge-1/2/0 {
    unit 0 {
        description B->D;
        family inet {
            address 172.16.1.1/24;
        }
    }
}
lo0 {
    unit 57 {
        family inet {
            address 10.0.0.1/32;
            address 10.0.0.2/32;
        }
    }
}

user@D# show protocols
bfd {
    traceoptions {
        file bfd-trace;
        flag all;
    }
}

user@B# show routing-options
static {
    route 192.168.47.0/24 {
        next-hop 172.16.1.2;
        bfd-liveness-detection {
            description Site- xxx;
            minimum-interval 1000;
        }
    }
}
```

Device D

```
user@D# show interfaces
ge-1/2/0 {
    unit 1 {
        description D->B;
        family inet {
            address 172.16.1.2/24;
        }
    }
}
```

Copyright © 2017, Juniper Networks, Inc.

```
lo0 {
  unit 2 {
    family inet {
      address 192.168.47.5/32;
      address 192.168.47.6/32;
    }
  }
}

user@D# show routing-options
static {
  route 0.0.0.0/0 {
    next-hop 172.16.1.1;
    bfd-liveness-detection {
      description Site - xxx;
      minimum-interval 1000;
    }
  }
}
```

Verification

Confirm that the configuration is working properly.

- Verifying That BFD Sessions Are Up on page 20
- Viewing Detailed BFD Events on page 21

Verifying That BFD Sessions Are Up

Purpose Verify that the BFD sessions are up, and view details about the BFD sessions.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@B> show bfd session extensive
                                                  Detect   Transmit
Address                 State     Interface       Time     Interval  Multiplier
172.16.1.2              Up        lt-1/2/0.0      3.000     1.000        3
 Client Static, description Site-xxx, TX interval 1.000, RX interval 1.000
 Session up time 00:14:30
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated, routing table index 172
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 2, remote discriminator 1
 Echo mode disabled/inactive

1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```


Copyright © 2017, Juniper Networks, Inc.

NOTE: The **description Site- <*xxx*>** is supported only on the SRX Series devices.

If each client has more than one description field, then it displays "and more" along with the first description field.

```
user@D> show bfd session extensive
                                                  Detect   Transmit
Address                 State     Interface       Time     Interval  Multiplier
172.16.1.1              Up        lt-1/2/0.1      3.000     1.000        3
 Client Static, TX interval 1.000, RX interval 1.000
 Session up time 00:14:35
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Replicated, routing table index 170
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 1, remote discriminator 2
 Echo mode disabled/inactive

1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

Meaning The **TX interval 1.000, RX interval 1.000** output represents the setting configured with the **minimum-interval** statement. All of the other output represents the default settings for BFD. To modify the default settings, include the optional statements under the **bfd-liveness-detection** statement.

Viewing Detailed BFD Events

Purpose View the contents of the BFD trace file to assist in troubleshooting, if needed.

Action From operational mode, enter the **file show /var/log/bfd-trace** command.

```
user@B> file show /var/log/bfd-trace
Nov 23 14:26:55    Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
 78 6d 69 74 20 72
Nov 23 14:26:55 PPM Trace: BFD periodic xmit rt tbl index 172
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 108:
Nov 23 14:26:55    IfIndex (3) len 4: 0
Nov 23 14:26:55    Protocol (1) len 1: BFD
Nov 23 14:26:55    Data (9) len 83: (hex) 70 70 6d 64 5f 62 66 64 5f 73 65 6e 64
 6d 73 67 20 3a 20
Nov 23 14:26:55 PPM Trace: ppmd_bfd_sendmsg : socket 12 len 24, ifl 78 src
172.16.1.1 dst 172.16.1.2 errno 65
Nov 23 14:26:55 Received Downstream TraceMsg (22) len 93:
Nov 23 14:26:55    IfIndex (3) len 4: 0
Nov 23 14:26:55    Protocol (1) len 1: BFD
Nov 23 14:26:55    Data (9) len 68: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
 78 6d 69 74 20 74
```

Meaning BFD messages are being written to the trace file.

Copyright © 2017, Juniper Networks, Inc.

Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection

This example shows how to configure a static route with multiple possible next hops. Each next hop has Bidirectional Forwarding Detection (BFD) enabled.

- Requirements on page 22
- Overview on page 22
- Configuration on page 23
- Verification on page 25

Requirements

In this example, no special configuration beyond device initialization is required.

Overview

In this example, Device B has the static route **192.168.47.0/24** with two possible next hops. The two next hops are defined using two **qualified-next-hop** statements. Each next hop has BFD enabled.

BFD is also enabled on Device D because BFD must be enabled on both ends of the connection.

A next hop is included in the routing table if the BFD session is up. The next hop is removed from the routing table if the BFD session is down.

See Figure 2 on page 22.

Figure 2: BFD Enabled on Qualified Next Hops

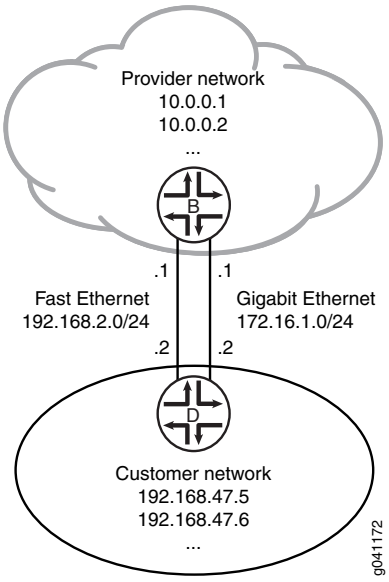

Copyright © 2017, Juniper Networks, Inc.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device B

```
set interfaces fe-0/1/0 unit 2 description secondary-B->D
set interfaces fe-0/1/0 unit 2 family inet address 192.168.2.1/24
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set routing-options static route 192.168.47.0/24 qualified-next-hop 192.168.2.2
  bfd-liveness-detection minimum-interval 60
set routing-options static route 192.168.47.0/24 qualified-next-hop 172.16.1.2
  bfd-liveness-detection minimum-interval 60
```

Device D

```
set interfaces fe-0/1/0 unit 3 description secondary-D->B
set interfaces fe-0/1/0 unit 3 family inet address 192.168.2.2/24
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set routing-options static route 0.0.0.0/0 qualified-next-hop 192.168.2.1
set routing-options static route 0.0.0.0/0 qualified-next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 60
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a static route with two possible next hops, both with BFD enabled:

1. On Device B, configure the interfaces.

   ```
   [edit interfaces fe-0/1/0]
   user@B# set unit 2 description secondary-B->D
   user@B# set unit 2 family inet address 192.168.2.1/24

   [edit interfaces ge-1/2/0]
   user@B# set unit 0 description B->D
   user@B# set unit 0 family inet address 172.16.1.1/24
   ```

2. On Device B, configure the static route with two next hops, both with BFD enabled.

   ```
   [edit routing-options static route 192.168.47.0/24]
   user@B# set qualified-next-hop 192.168.2.2 bfd-liveness-detection minimum-interval
     60
   user@B# set qualified-next-hop 172.16.1.2 bfd-liveness-detection minimum-interval
     60
   ```

3. On Device D, configure the interfaces.

   ```
   [edit interfaces fe-0/1/0]
   user@D# set unit 3 description secondary-D->B
   user@D# set unit 3 family inet address 192.168.2.2/24

   [edit interfaces ge-1/2/0]
   user@D# set unit 1 description D->B
   ```

Copyright © 2017, Juniper Networks, Inc.

```
user@D# set unit 1 family inet address 172.16.1.2/24
```

4. On Device D, configure a BFD-enabled default static route with two next hops to the provider network.

   In this case, BFD is enabled on the route, not on the next hops.

```
[edit routing-options static route 0.0.0.0/0]
user@D# set qualified-next-hop 192.168.2.1
user@D# set qualified-next-hop 172.16.1.1
user@D# set bfd-liveness-detection minimum-interval 60
```

Results Confirm your configuration by issuing the **show interfaces** and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@B# show interfaces
fe-0/1/0 {
    unit 2 {
        description secondary-B->D;
        family inet {
            address 192.168.2.1/24;
        }
    }
}
ge-1/2/0 {
    unit 0 {
        description B->D;
        family inet {
            address 172.16.1.1/24;
        }
    }
}
```

```
user@B# show routing-options
static {
    route 192.168.47.0/24 {
        qualified-next-hop 192.168.2.2 {
            bfd-liveness-detection {
                minimum-interval 60;
            }
        }
        qualified-next-hop 172.16.1.2 {
            bfd-liveness-detection {
                minimum-interval 60;
            }
        }
    }
}
```

```
user@D# show interfaces
fe-0/1/0 {
    unit 3 {
        description secondary-D->B;
        family inet {
            address 192.168.2.2/24;
        }
```


Copyright © 2017, Juniper Networks, Inc.

```
    }
  }
  ge-1/2/0 {
    unit 1 {
      description D->B;
      family inet {
        address 172.16.1.2/24;
      }
    }
  }
```

```
user@D# show routing-options
static {
  route 0.0.0.0/0 {
    qualified-next-hop 192.168.2.1;
    qualified-next-hop 172.16.1.1;
    bfd-liveness-detection {
      minimum-interval 60;
    }
  }
}
```

If you are done configuring the devices, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Checking the Routing Tables on page 25
- Verifying the BFD Sessions on page 26
- Removing BFD from Device D on page 26
- Removing BFD from One Next Hop on page 27

Checking the Routing Tables

Purpose Make sure that the static route appears in the routing table on Device B with two possible next hops.

Copyright © 2017, Juniper Networks, Inc.

Action

```
user@B> show route 192.168.47.0 extensive
inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
192.168.47.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.47.0/24 -> {192.168.2.2}
        *Static Preference: 5
                Next hop type: Router
                Address: 0x9334010
                Next-hop reference count: 1
                Next hop: 172.16.1.2 via ge-1/2/0.0
                Next hop: 192.168.2.2 via fe-0/1/0.2, selected
                State: <Active Int Ext>
                Age: 9
                Task: RT
                Announcement bits (1): 3-KRT
                AS path: I
```

Meaning Both next hops are listed. The next hop 192.168.2.2 is the selected route.

Verifying the BFD Sessions

Purpose Make sure that the BFD sessions are up.

Action

```
user@B> show bfd session

                                                Detect   Transmit
Address          State     Interface      Time     Interval  Multiplier
172.16.1.2       Up        ge-1/2/0.0     0.720     0.240        3
192.168.2.2      Up        fe-0/1/0.2     0.720     0.240        3

2 sessions, 2 clients
Cumulative transmit rate 8.3 pps, cumulative receive rate 8.3 pps
```

Meaning The output shows that the BFD sessions are up.

Removing BFD from Device D

Purpose Demonstrate what happens when the BFD session is down for both next hops.

Action

1. Deactivate BFD on Device D.

   ```
   [edit routing-options static route 0.0.0.0/0]
   user@D# deactivate bfd-liveness-detection
   user@D# commit
   ```

2. Rerun the **show bfd session** command on Device B.

   ```
   user@B> show bfd session

                                                  Detect   Transmit
   Address                State     Interface     Time     Interval  Multiplier
   172.16.1.2              Down      ge-1/2/0.0     3.000     1.000        3
   192.168.2.2             Down      fe-0/1/0.2     3.000     1.000        3

   2 sessions, 2 clients
   Cumulative transmit rate 2.0 pps, cumulative receive rate 2.0 pps
   ```


Copyright © 2017, Juniper Networks, Inc.

3. Rerun the **show route 192.168.47.0** command on Device B.

```
user@B> show route 192.168.47.0
```

Meaning As expected, when the BFD sessions are down, the static route is removed from the routing table.

Removing BFD from One Next Hop

Purpose Demonstrate what happens when only one next hop has BFD enabled.

Action

1. If it is not already deactivated, deactivate BFD on Device D.

```
[edit routing-options static route 0.0.0.0/0]
user@D# deactivate bfd-liveness-detection
user@D# commit
```

2. Deactivate BFD on one of the next hops on Device B.

```
[edit routing-options static route 192.168.47.0/24 qualified-next-hop 172.16.1.2]
user@B# deactivate bfd-liveness-detection
user@B# commit
```

3. Rerun the **show bfd session** command on Device B.

```
user@B> show bfd session

                                       Detect   Transmit
Address      State     Interface      Time     Interval  Multiplier
192.168.2.2   Down      fe-0/1/0.2     3.000     1.000        3
```

4. Rerun the **show route 192.168.47.0 extensive** command on Device B.

```
user@B> show route 192.168.47.0 extensive

inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
192.168.47.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.47.0/24 -> {172.16.1.2}
        *Static Preference: 5
                Next hop type: Router, Next hop index: 624
                Address: 0x92f0178
                Next-hop reference count: 3
                Next hop: 172.16.1.2 via ge-1/2/0.0, selected
                State: <Active Int Ext>
                Age: 2:36
                Task: RT
                Announcement bits (1): 3-KRT
                AS path: I
```

Meaning As expected, the BFD session is down for the 192.168.2.2 next hop. The 172.16.1.2 next hop remains in the routing table, and the route remains active, because BFD is not a condition for this next hop to remain valid.

Copyright © 2017, Juniper Networks, Inc.

Release History Table

Release	Description
16.1R1	Starting in Junos OS Release 16.1R1, the inline BFD sessions are supported on integrated routing and bridging (IRB) interfaces.
15.1F6	Inline BFD is supported on PTX3000 routers with third-generation FPCs starting in Junos OS Release 15.1F6 and 16.1R2.
15.1F3	Inline BFD is supported on PTX5000 routers with third-generation FPCs starting in Junos OS Release 15.1F3 and 16.1R2.
13.3	Starting in Junos OS Release 13.3, inline BFD is supported only on static MX Series routers with MPCs/MICs that have configured **enhanced-ip**.

Related Documentation

- Example: Configuring BFD Authentication for Static Routes on page 28
- *Example: Configuring BFD for OSPF*
- *Example: Configuring BFD for BGP*
- *Example: Configuring BFD for IS-IS*
- *Configuring PIM and the Bidirectional Forwarding Detection (BFD) Protocol*

Example: Configuring BFD Authentication for Static Routes

- Understanding BFD Authentication for Static Route Security on page 28
- Example: Configuring BFD Authentication for Securing Static Routes on page 30

Understanding BFD Authentication for Static Route Security

Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant.

NOTE: We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels.

Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over IPv4 and IPv6 static routes. BFD authentication is not supported on MPLS OAM sessions. BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

NOTE: EX3300 supports BFD over static routes only.

Copyright © 2017, Juniper Networks, Inc.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and level of authentication that can be configured:

- BFD Authentication Algorithms on page 29
- Security Authentication Keychains on page 30
- Strict Versus Loose Authentication on page 30

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.

NOTE: Nonstop active routing (NSR) is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

Copyright © 2017, Juniper Networks, Inc.

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled, and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Example: Configuring BFD Authentication for Securing Static Routes

This example shows how to configure Bidirectional Forwarding Detection (BFD) authentication for static routes.

- Requirements on page 30
- Overview on page 30
- Configuration on page 31
- Verification on page 34

Requirements

Junos OS Release 9.6 or later (Canda and United States version).

BFD authentication is only supported in the Canada and United States version of the Junos OS image and is not available in the export version.

Overview

You can configure authentication for BFD sessions running over IPv4 and IPv6 static routes. Routing instances and logical systems are also supported.

The following steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the static route.
2. Associate the authentication keychain with the static route.

Copyright © 2017, Juniper Networks, Inc.

3. Configure the related security authentication keychain. This must be configured on the main router.

TIP: We recommend that you specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit]
user@host> set routing-options static route ipv4 bfd-liveness-detection
  authentication loose-check
```

Figure 3 on page 31 shows the sample network.

Figure 3: Customer Routes Connected to a Service Provider

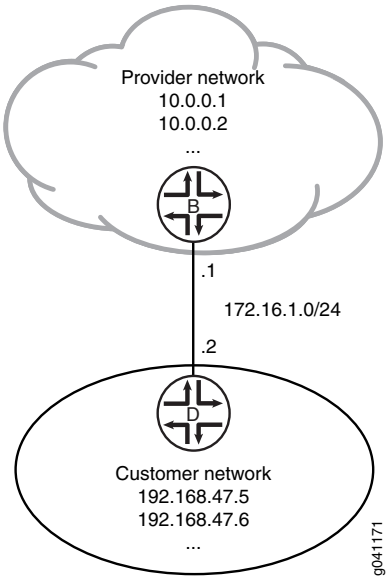

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device B

```
set interfaces ge-1/2/0 unit 0 description B->D
set interfaces ge-1/2/0 unit 0 family inet address 172.16.1.1/24
set interfaces lo0 unit 57 family inet address 10.0.0.1/32
set interfaces lo0 unit 57 family inet address 10.0.0.2/32
set routing-options static route 192.168.47.0/24 next-hop 172.16.1.2
set routing-options static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
  1000
set routing-options static route 192.168.47.0/24 bfd-liveness-detection description
  Site-xxx
set routing-options static route 192.168.47.0/24 bfd-liveness-detection authentication
  key-chain bfd-kc4
```

Copyright © 2017, Juniper Networks, Inc.

```
set routing-options static route 192.168.47.0/24 bfd-liveness-detection authentication
  algorithm keyed-sha-1
set security authentication-key-chains key-chain bfd-kc4 key 5 secret
  "$ABC123$ABC123$ABC123"
set security authentication-key-chains key-chain bfd-kc4 key 5 start-time
  "2011-1-1.12:00:00 -0800"
```

Device D

```
set interfaces ge-1/2/0 unit 1 description D->B
set interfaces ge-1/2/0 unit 1 family inet address 172.16.1.2/24
set interfaces lo0 unit 2 family inet address 192.168.47.5/32
set interfaces lo0 unit 2 family inet address 192.168.47.6/32
set routing-options static route 0.0.0.0/0 next-hop 172.16.1.1
set routing-options static route 0.0.0.0/0 bfd-liveness-detection minimum-interval 1000
set routing-options static route 0.0.0.0/0 bfd-liveness-detection authentication key-chain
  bfd-kc4
set routing-options static route 0.0.0.0/0 bfd-liveness-detection authentication algorithm
  keyed-sha-1
set security authentication-key-chains key-chain bfd-kc4 key 5 secret
  "$ABC123$ABC123$ABC123"
set security authentication-key-chains key-chain bfd-kc4 key 5 start-time
  "2011-1-1.12:00:00 -0800"
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure BFD for static routes:

1. On Device B, configure the interfaces.

```
[edit interfaces]
user@B# set ge-1/2/0 unit 0 description B->D
user@B# set ge-1/2/0 unit 0 family inet address 172.16.1.1/24


user@B# set lo0 unit 57 family inet address 10.0.0.1/32
user@B# set lo0 unit 57 family inet address 10.0.0.2/32
```

2. On Device B, create a static route and set the next-hop address.

```
[edit routing-options]
user@B# set static route 192.168.47.0/24 next-hop 172.16.1.2
```

3. On Device B, configure BFD for the static route.

```
[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection minimum-interval
  1000
set routing-options static route 192.168.47.0/24 bfd-liveness-detection description
  Site-xxx
```

4. On Device B, specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use for BFD authentication on the static route.

```
[edit routing-options]
user@B# set static route 192.168.47.0/24 bfd-liveness-detection authentication
  algorithm keyed-sha-1
```


Copyright © 2017, Juniper Networks, Inc.

NOTE: Nonstop active routing (NSR) is not supported with the meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

5. On Device B, specify the keychain to be used to associate BFD sessions on the specified route with the unique security authentication keychain attributes.

   This should match the keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

   ```
   [edit routing-options]
   user@B# set static route 192.168.47.0/24 bfd-liveness-detection authentication
     key-chain bfd-kc4
   ```

6. On Device B, specify the unique security authentication information for BFD sessions:

   - The matching keychain name as specified in Step 5.
   - At least one key, a unique integer between **0** and **63**. Creating multiple keys allows multiple clients to use the BFD session.
   - The secret data used to allow access to the session.
   - The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

   ```
   [edit security authentication-key-chains key-chain bfd-kc4]
   user@B# set key 5 secret "$ABC123$ABC123$ABC123"
   user@B# set key 5 start-time "2011-1-1.12:00:00 -0800"
   ```

7. If you are done configuring Device B, commit the configuration.

   ```
   [edit]
   user@B# commit
   ```

8. Repeat the configuration on Device D.

   The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

Results

Confirm your configuration by issuing the **show interfaces**, **show routing-options**, and **show security** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

Device B

```
user@B# show interfaces
ge-1/2/0 {
  unit 0 {
    description B->D;
    family inet {
      address 172.16.1.1/24;
    }
```

Copyright © 2017, Juniper Networks, Inc.

```
    }
  }
  lo0 {
    unit 57 {
      family inet {
        address 10.0.0.1/32;
        address 10.0.0.2/32;
      }
    }
  }

  user@B# show routing-options
  static {
    route 192.168.47.0/24 {
      next-hop 172.16.1.2;
      bfd-liveness-detection {
        description Site- xxx;
        minimum-interval 1000;
        authentication {
          key-chain bfd-kc4;
          algorithm keyed-sha-1;
        }
      }
    }
  }

  user@B# show security
  authentication-key-chains {
    key-chain bfd-kc4 {
      key 5 {
        secret "$ABC123$ABC123$ABC123"; ## SECRET-DATA
        start-time "2011-1-1.12:00:00 -0800";
      }
    }
  }
```

Verification

Confirm that the configuration is working properly.

- Verifying That BFD Sessions Are Up on page 34
- Viewing Details About the BFD Session on page 35
- Viewing Extensive BFD Session Information on page 35

Verifying That BFD Sessions Are Up

Purpose Verify that the BFD sessions are up.

Action From operational mode, enter the **show bfd session** command.

```
user@B> show bfd session
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
172.16.1.2               Up        ge-1/2/0.0     3.000     1.000        3
```


Copyright © 2017, Juniper Networks, Inc.

```
1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

Meaning The command output shows that the BFD session is up.

Viewing Details About the BFD Session

Purpose View details about the BFD sessions and make sure that authentication is configured.

Action From operational mode, enter the **show bfd session detail** command.

```
user@B> show bfd session detail
                                                Detect   Transmit
Address                 State     Interface     Time     Interval  Multiplier
172.16.1.2              Up        ge-1/2/0.0    3.000     1.000        3
 Client Static, TX interval 1.000, RX interval 1.000, Authenticate
 Session up time 00:53:58
 Local diagnostic NbrSignal, remote diagnostic None
 Remote state Up, version 1
 Logical system 9, routing table index 22

1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

Meaning In the command output, **Authenticate** is displayed to indicate that BFD authentication is configured.

Viewing Extensive BFD Session Information

Purpose View more detailed information about the BFD sessions.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@B> show bfd session extensive
Address                 State     Interface     Time     Interval  Multiplier
172.16.1.2              Up        ge-1/2/0.0    3.000     1.000        3
 Client Static, description Site-xxx, TX interval 1.000, RX interval 1.000,
Authenticate
        keychain bfd-kc4, algo keyed-sha-1, mode strict
 Session up time 01:39:45
 Local diagnostic NbrSignal, remote diagnostic None
 Remote state Up, version 1
 Logical system 9, routing table index 22
 Min async interval 1.000, min slow interval 1.000
 Adaptive async TX interval 1.000, RX interval 1.000
 Local min TX interval 1.000, minimum RX interval 1.000, multiplier 3
 Remote min TX interval 1.000, min RX interval 1.000, multiplier 3
 Local discriminator 3, remote discriminator 4
 Echo mode disabled/inactive
 Authentication enabled/active, keychain bfd-kc4, algo keyed-sha-1, mode strict

1 sessions, 1 clients
Cumulative transmit rate 1.0 pps, cumulative receive rate 1.0 pps
```

Meaning In the command output, **Authenticate** is displayed to indicate that BFD authentication is configured. The output for the **extensive** command provides the keychain name, the authentication algorithm, and the mode for each client in the session.

Copyright © 2017, Juniper Networks, Inc.

NOTE: The **description Site- <*xxx*>** is supported only on the SRX Series devices.

If each client has more than one description field, then it displays "and more" along with the first description field.

Related Documentation

- Examples: Configuring BFD for Static Routes on page 11

Understanding the Unified Forwarding Table

- Using the Unified Forwarding Table to Optimize Address Storage on page 36
- MAC Address and Host Address Memory Allocation on page 36
- LPM Table Memory Allocation on page 38

Using the Unified Forwarding Table to Optimize Address Storage

On QFX5100, EX4600, and OCX1100 switches, you can control the allocation of forwarding table memory available to store the following:

- MAC addresses.
- Layer 3 host entries.
- Longest prefix match (LPM) table entries.

NOTE: Starting with Junos OS 13.2X51-D15, you can allocate more memory to store prefixes in the range /65 to /127 range.

NOTE: Platform support depends on the Junos OS release in your installation.

This feature gives you the flexibility to configure your switch to match the needs of your particular network environment. For example, you might configure the switch to store more MAC addresses in a Layer 2 network, such as a virtualized network with many servers and virtualized machines. On the other hand, if your switch is located in the routing core of a network or participates in an IP fabric, you probably want to maximize the number of routing table entries it can store. In this case, you would configure it to use the **lpm-profile**, which provides the most longest prefix match table entries.

MAC Address and Host Address Memory Allocation

There are several profiles that allocate memory differently for MAC addresses and host addresses. You configure the mix that best meets your needs by choosing the appropriate profile. Table 3 on page 37 lists the profiles you can choose and the associated maximum

Copyright © 2017, Juniper Networks, Inc.

values for the MAC address and host table entries on QFX3500, QFX3600, and QFX5100 switches.

Table 3: Unified Forwarding Table Profiles on QFX3500, QFX3600, and QFX5100 Switches

Profile Name	MAC Table	Host Table (unicast and multicast addresses)					
	MAC Addresses	IPv4 unicast	IPv6 unicast	IPv4 (*, G)	IPv4 (S, G)	IPv6 (*, G)	IPv6 (S, G)
l2-profile-one	288K	16K	8K	8K	8K	4K	4K
l2-profile-two	224K	80K	40K	40K	40K	20K	20K
l2-profile-three (default)	160K	144K	72K	72K	72K	36K	36K
l3-profile	96K	208K	104K	104K	104K	52K	52K
lpm-profile	32K	16K	8K	8K	8K	4K	4K
lpm-profile with **unicast-in-lpm** option	32K	(stored in LPM table)	(stored in LPM table)	8K	8K	4K	4K

NOTE: On QFX5100 and QFX5200 switches, IPv4 and IPv6 host routes with ECMP next hops are stored in the host table. On QFX3500 and QFX3600 switches, these routes are stored in the LPM table.

Table 4 on page 37 lists the profiles you can choose and the associated maximum values for the MAC address and host table entries on QFX5200 switches.

Table 4: Unified Forwarding Table Profiles on QFX5200 Switches

Profile Name	MAC Table	Host Table (unicast and multicast addresses)						
	MAC Addresses	IPv4 unicast	IPv6 unicast	IPv4 (*, G)	IPv4 (S, G)	IPv6 (*, G)	IPv6 (S, G)	Exact-Match
l2-profile-one	136K	8K	4K	4K	4K	2K	2K	0
l2-profile-two	104K	40K	20K	20K	20K	10K	10K	0
l2-profile-three (default)	72K	72K	36K	36K	36K	18K	18K	0
l3-profile	40K	104K	52K	52K	52K	26K	26K	0
exact-match-profile	8K	8K	4K	4K	4K	2K	2K	64K
lpm-profile	8K	8K	4K	4K	4K	2K	2K	0

Copyright © 2017, Juniper Networks, Inc.

Table 4: Unified Forwarding Table Profiles on QFX5200 Switches *(continued)*

Profile Name	MAC Table	Host Table (unicast and multicast addresses)						
lpm-profile with **unicast-in-lpm** option	?	(stored in LPM table)	(stored in LPM table)	?	?	?	?	0

Note that all entries in the host table share the same memory space. If the host table stores the maximum number of entries for any given type, the entire shared table is full and is unable to accommodate *any* entries of any other type. As you can see, different entry types occupy different amounts of memory. For example, an IPv6 unicast address occupies twice as much memory as an IPv4 unicast address, and an IPv6 multicast address occupies four times as much memory as an IPv4 unicast address.
Table 5 on page 38 lists various valid combinations that the host table can store if you use the **l2-profile-one** profile on a QFX5100 switch. Each row in the table represents a case in which the host table is full and cannot accommodate any more entries. .

Table 5: Example Host Table Combinations Using l2-profile-one on a QFX5100 Switch

IPv4 unicast	IPv6 unicast	IPv4 multicast (*, G)	IPv4 multicast (S, G)	IPv6 multicast (*, G)	IPv6 multicast (S, G)
16K	0	0	0	0	0
12K	2K	0	0	0	0
12K	0	2K	2K	0	0
8K	4K	0	0	0	0
4K	2K	2K	2K	0	0
0	4K	0	0	1K	1K

LPM Table Memory Allocation

You configure the memory allocation for LPM table entries differently depending on which version of Junos OS you use. To learn how to configure memory allocation for LPM table entries see "Configuring the Unified Forwarding Table" on page 39. Starting with Junos OS 14.1X53-D30 you can free memory in the host table by using the **unicast-in-lpm** option with the **lpm-profile** to store IPv4 and IPv6 unicast addresses in the LPM table instead of the host table. See "Configuring the lpm-profile With Junos OS Release 14.1x53-D30 and Later" on page 45.

Copyright © 2017, Juniper Networks, Inc.

Release History Table

Release	Description
14.1X53-D30	Starting with Junos OS 14.1X53-D30 you can free memory in the host table by using the **unicast-in-lpm** option with the **lpm-profile** to store IPv4 and IPv6 unicast addresses in the LPM table instead of the host table.
13.2X51-D15	Starting with Junos OS 13.2X51-D15, you can allocate more memory to store prefixes in the range /65 to /127 range.

Related Documentation

- Configuring the Unified Forwarding Table on page 39

Configuring the Unified Forwarding Table

Traditionally, forwarding tables have been statically defined and have supported only a fixed number of entries for each type of address stored in the tables. The Unified Forwarding Table feature lets you optimize how your switch allocates forwarding-table memory for different types of addresses. You can choose one of five unified forwarding table profiles. Each profile allocates a different maximum amount of memory for Layer 2, Layer 3 host, and longest prefix match (LPM) entries. In addition to selecting a profile, you can also select how much additional memory to allocate for LPM entries.

Two profiles allocate higher percentages of memory to Layer 2 addresses. A third profile allocates a higher percentage of memory to Layer 3 host address, while a fourth profile allocates a higher percentage of memory to LPM entries. There is a default profile configured that allocates an equal amount of memory to Layer 2 and Layer 3 host addresses with the remainder allocated to LPM entries. For a switch in a virtualized network that handles a great deal of Layer 2 traffic, you would choose a profile that allocates a higher percentage of memory to Layer 2 addresses. For a switch that operates in the core of the network, you would choose a profile that allocates a higher percentage of memory to LPM entries.

On QFX5200 Series switches only, you can also configure a custom profile that allows you to partition shared memory banks among the different types of forwarding table entries. These shared memory banks have a total memory equal to 128,000 IPv4 unicast addresses. For more information about configuring the custom profile, see *Example: Configuring a Unified Forwarding Table Custom Profile on QFX Series Switches*.

- Configuring a Unified Forwarding Table Profile on page 40
- Configuring the Memory Allocation for Longest Prefix Match Entries on page 41

Copyright © 2017, Juniper Networks, Inc.

Configuring a Unified Forwarding Table Profile

To configure a unified forwarding table profile:

Specify a forwarding-table profile.

```
[edit chassis forwarding-options]
user@switch# set profile-name
```

For example, to specify the profile that allocates the highest percentage of memory to Layer 2 traffic:

```
[edit chassis forwaridng-options]
user@switch# set l2-profile-one
```

CAUTION: When you configure and commit a profile, the Packet Forwarding Engine automatically restarts and all the data interfaces on the switch go down and come back up (the management interfaces are unaffected).

However, starting with Junos OS Release 14.1X53-D40, 15.1R5, and 16.1R3, Packet Forwarding Engines on switches in a Virtual Chassis or Virtual Chassis Fabric (VCF) do not automatically restart upon configuring a unified forwarding table profile change. This behavior avoids Virtual Chassis or VCF instability after the change propagates to member switches and multiple Packet Forwarding Engines automatically restart at the same time. Instead, a message is displayed at the CLI prompt and logged to the switch's system log to notify you that the profile change does not take effect until the next time you reboot the Virtual Chassis or VCF. We recommend that you plan to make profile changes only when you can perform a Virtual Chassis or VCF system reboot immediately after committing the configuration update. Otherwise, the Virtual Chassis or VCF could become inconsistent if one or more members have a problem and restart with the new configuration before a planned system reboot activates the change on all members.

NOTE: You can configure only one profile for the entire switch.

NOTE: The l2-profile-three is configured by default.

NOTE: If the host table stores the maximum number of entries for any given type, the entire table is full and is unable to accommodate *any* entries of any other type. Keep in mind that an IPv6 unicast address occupies twice as much memory as an IPv4 unicast address, and an IPv6 multicast address occupies four times as much memory as an IPv4 unicast address. For more information about valid combinations of table entries see *Understanding the Unified Forwarding Table*.

Copyright © 2017, Juniper Networks, Inc.

Configuring the Memory Allocation for Longest Prefix Match Entries

In addition to choosing a profile, you can further optimize memory allocation for longest prefix match (LPM) entries by configuring how many IPv6 prefixes to store with lengths from /65 through /127. The switch uses LPM entries during address lookup to match addresses to the most-specific (longest) applicable prefix. Prefixes of this type are stored in the space for ternary content addressable memory (TCAM). Changing the default parameters makes this space available for LPM entries. Increasing the amount of memory available for these IPv6 prefixes reduces by the same amount how much memory is available to store IPv4 unicast prefixes and IPv6 prefixes with lengths equal to or less than 64.

The procedures for configuring the LPM table are different, depending on which version of Junos OS you are using. In the initial releases that UFT is supported, Junos OS Relases 3.2x51-D10 and 13.2X52-10, you can only increase the amount of memory allocated to IPv6 prefixes with lengths from /65 through /127 for any profile, except for **lpm-profile**. Starting with Junos OS Release 13.2x51-D15, you can also allocate either less or no memory for IPv6 prefixes with lengths in the range /65 through /127, depending on which profile is configured. For the **lpm-profie**, however, the only change you can make to the default parameters is to allocate no memory for these types of prefixes.

- Configuring the LPM Table With Junos OS Releases 13.2X51-D10 and 13.2X52-D10 on page 41
- Configuring the LPM Table With Junos OS Release 13.2x51-D15 and Later on page 42

Configuring the LPM Table With Junos OS Releases 13.2X51-D10 and 13.2X52-D10

In Junos OS Releases 13.2x51-D10 and 13.2X52-D10, by default, the switch allocates memory for 16 IPv6 with prefixes with lengths in the range /65 through /127. You can configure the switch to allocate more memory for IPv6 prefixes with lengths in the range /65 through /127.

To allocate more memory for IPv6 prefixes in the range /65 through /127:

1. Choose a forwarding table profile.

   ```
   [edit chassis forwarding-options]
   user@swtitch# set profile-name
   ```

   For example, to specify the profile that allocates the highest percentage of memory to Layer 2 traffic:

   ```
   [edit chassis forwarding-options]
   user@swtitch# set l2-profile-one
   ```

2. Select how much memory to allocate for IPv6 prefixes in the range /65 thorugh 127.

   ```
   [edit chassis forwarding-options profile-name]
   user@swtitch# set num-65-127-prefix number
   ```

   For example, to specify to allocate memory for 32 IPv6 prefixes in the range /65 through 127:

   ```
   [edit chassis forwarding-options l2-profile-one]
   user@switch# set num-65-127-prefix 2
   ```

Copyright © 2017, Juniper Networks, Inc.

NOTE: When you configure and commit the **num-65-127-prefix** *number* statement, all the data interfaces on the switch restart. The management interfaces are unaffected.

The **num-65-127-prefix** *number* statement is not supported on the **lpm-profile**.

Table 6 on page 42 provides examples of valid combinations that the LPM table can store for all profiles except for the **lpm-profile** and **custom-profile**. Once again, each row in the table represents a case in which the table is full and cannot accommodate any more entries.

Table 6: Example LPM Table Combinations Using L2 and L3 profiles With Junos OS 13.2X51-D10 and 13.2X52-D10

num-65-127-prefix Value	IPv4 Entries	IPv6 Entries (Prefix <= 64)	IPv6 Entries (Prefix >= 65)
1 (default)	16K-16	0K	16
1 (default)	0K	8K-16	16
1 (default)	8K-16	4K	16
64	4K	4K	1K
64	2K	5K	1K
64	0K	6K	1K
128	4K	2K	2K
128	2K	3K	2K
128	0K	4K	2K

Configuring the LPM Table With Junos OS Release 13.2x51-D15 and Later

- Configuring Layer 2 and Layer 3 Profiles With Junos OS Release 13.2x51-D15 or Later on page 42
- Configuring the lpm-profile With Junos OS Release 13.2x51-D15 and Later on page 44
- Configuring the lpm-profile With Junos OS Release 14.1x53-D30 and Later on page 45

Configuring Layer 2 and Layer 3 Profiles With Junos OS Release 13.2x51-D15 or Later

Starting in Junos OS Release 13.2x51-D15 , you can configure the switch to allocate forwarding table memory for as many as 4,000 IPv6 prefixes with lengths in the range /65 through /127 for any profile other than the **lpm-profile** or **custom-profile**. You can also specify to allocate no memory for these IPv6 entries. The default is 1,000 entries for IPv6 prefixes with lengths in the range /65 through /127. Previously, the maximum you could

Copyright © 2017, Juniper Networks, Inc.

configure was for 2,048 entries for IPv6 prefixes with lengths in the range /65 through /127. The minimum number of entries was previously 16, which was the default.

To specify how much forwarding table memory to allocate for IPv6 prefixes with length in the range /65 through /127:

1. Choose a forwarding table profile.

   ```
   [edit chassis forwarding-options]
   user@swtitch# set profile-name
   ```

   For example, to specify the profile that allocates the highest percentage of memory to Layer 2 traffic:

   ```
   [edit chassis forwarding-options]
   user@swtitch# set l2-profile-one
   ```

2. Select how much memory to allocate for IPv6 prefixes in the range /65 thorugh 127.

   ```
   [edit chassis forwarding-options profile-name]
   user@swtitch# set num-65-127-prefix number
   ```

   For example, to specify to allocate memory for 2,000 IPv6 prefixes in the range /65 through 127:

   ```
   [edit chassis forwarding-options l2-profile-one]
   user@switch# set num-65-127-prefix 2
   ```

Table 7 on page 43 shows the numbers of entries that you can allocate by using the **num-65-127-prefix** statement starting with Junos OS 13.2X51-D15. Each row represents a case in which the table is full and cannot accommodate any more entries.

Table 7: LPM Table Combinations for L2 and L3 profiles With Junos OS 13.2X51-D15 and Later

num-65-127-prefix Value	IPv4 Entries	IPv6 Entries (Prefix <= 64)	IPv6 Entries (Prefix >= 65)
0	16K	8K	0K
1 (default)	12K	6K	1K
2	8K	4K	2K
3	4K	2K	3K
4	0K	0K	4K

CAUTION: When you configure and commit a profile change with the **num-65-127-prefix** *number* statement, the Packet Forwarding Engine automatically restarts and all the data interfaces on the switch go down and come back up (the management interfaces are unaffected).

However, starting with Junos OS Release 14.1X53-D40, 15.1R5, and 16.1R3, Packet Forwarding Engines on switches in a Virtual Chassis or Virtual Chassis Fabric (VCF) do not automatically restart upon configuring a unified

Copyright © 2017, Juniper Networks, Inc.

forwarding table profile change. This behavior avoids Virtual Chassis or VCF instability when the change propagates to member switches and multiple Packet Forwarding Engines restart at the same time. Instead, a message is displayed at the CLI prompt and logged to the switch's system log to notify you that the profile change does not take effect until the next time you reboot the Virtual Chassis or VCF. We recommend that you plan to make profile changes only when you can perform a Virtual Chassis or VCF system reboot immediately after committing the configuration update. Otherwise, the Virtual Chassis or VCF could become inconsistent if one or more members have a problem and restart with the new configuration before a planned system reboot activates the change on all members.

Configuring the lpm-profile With Junos OS Release 13.2x51-D15 and Later

Starting with Junos OS release 13.2x51-D15 you can configure the **lpm-profile** profile not to allocate any memory for IPv6 entries with prefix lengths from /65 through /127. These are the default maximum values allocated for LPM memory for the **lpm-profile** by address type:

- 128K of IPv4 prefixes
- 16K of IPv6 prefixes (all lengths)

NOTE: The memory allocated for each address type represents the maximum default value for all LPM memory.

To configure the **lpm-profile** not to allocate forwarding-table memory for IPv6 entries with prefixes from /65 through /127, thus allocating more memory for IPv4:

Specify to disable forwarding-table memory for IPv6 prefixes with lengths in the range /65 through /127.

```
[edit chassis forwarding-options lpm-profile]
user@switch# set prefix-65-127-disable
```

For example, on the QFX3500, QFX3600, and QFX5100 switches only, if you use the **prefix-65-127-disable** option, each of the following combinations are valid:

- 100K IPv4 and 28K IPv6 /64 or shorter prefixes.
- 64K IPv4 and 64K IPv6 /64 or shorter prefixes.
- 128K IPv4 and 0K IPv6 /64 or shorter prefixes.
- 0K IPv4 and 128K IPv6 /64 or shorter prefixes.

NOTE: On the QFX5200 switches, when you configure the **prefix-65-127-disable** statement, the maximum number of IPv6 entries with prefixes equal to or shorter than 64 is 98,000.

Copyright © 2017, Juniper Networks, Inc.

Configuring the lpm-profile With Junos OS Release 14.1x53-D30 and Later

Starting in Junos OS Release 15.1x53-D30, you can configure the **lpm-profile** profile to store unicast IPv4 and IPv6 host addresses in the LPM table , thereby freeing memory in the host table. Unicast IPv4 and IPv6 addresses are stored in the LPM table instead of the host table, as shown in Table 8 on page 45 for QFX3500, QFX3600, and QFX5100 switches. You can use this option in conjunction with the option to allocate no memory in the LPM table for IPv6 entries with prefix lengths in the range /65 through /127. Together, these options maximize the amount of memory available for IPv4 unicast entries and IPv6 entries with prefix lengths equal to or less than 64.

Table 8: lpm-profile with unicast-in-lpm Option for QFX3500, QFX3600, and QFX5100 Switches

prefix-65-127-disable	MAC Table	Host Table (multicast addresses)						LPM Table unicast addresses)		
	MAC	IPv4 unicast	IPv6 unicast	IPv4 (*, G)	IPv4 (S, G)	IPv6 (*, G)	IPv6 (S, G)	IPv4 unicast	IPv6 unicast (</65)	IPv6 unicast (>/64)
No	32K	0	0	8K	8K	4K	4K	128K	16K	16K
Yes	32K	0	0	8K	8K	4K	4K	128K	128K	0

Note that all entries in each table share the same memory space. If a table stores the maximum number of entries for any given type, the entire shared table is full and is unable to accommodate any entries of any other type. For example, if you use the the **unicast-in-lpm** option and there are 128K IPv4 unicast addresses stored in the LPM table, the entire LPM table is full and no IPv6 addresses can be stored. Similarly, if you use the **unicast-in-lpm** option but do not use the **prefix-65-127-disable** option, and 16K IPv6 addresses with prefixes shorter than /65 are stored, the entire LPM table is full and no additional addresses (IPv4 or IPv6) can be stored.

To configure the **lpm-profile** to store unicast IPv4 entries and IPv6 entries with prefix lengths equal to or less than 64 in the LPM table:

1. Specify the option to store these entries in the LPM table.

   ```
   [edit chassis forwarding-options lpm-profile]
   user@switch# set unicast-in-lpm
   ```

2. (Optional) Specify to allocate no memory for in the LPM table for IPv6 prefixes with length in the range /65 through /127:

   ```
   [edit chassis forwarding-options lpm-profile]
   user@switch# set prefix-65-127-disable
   ```

Related Documentation

- *Understanding the Unified Forwarding Table*

Monitoring Routing Information

Purpose Use the monitoring functionality to view the **inet.0** routing table on the routing device.

Copyright © 2017, Juniper Networks, Inc.

Action To view the routing table, enter the following commands in the CLI interface:

- **show route terse**
- **show route detail**

Meaning Table 9 on page 46 describes the different filters, their functions, and the associated actions.

Table 10 on page 46 summarizes key output fields in the routing information display.

Table 9: Filtering Route Messages

Field	Function	Your Action
Destination Address	Specifies the destination address of the route.	Enter the destination address.
Next hop address	Specifies the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.	Enter the next hop address.
Best route	Specifies only the best route available.	Select the view details of the best route.
Inactive routes	Specifies the inactive routes.	Select the view details of inactive routes.
Exact route	Specifies the exact route.	Select the view details of the exact route.
Hidden routes	Specifies the hidden routes.	Select the view details of hidden routes.
Search	Applies the specified filter and displays the matching messages.	To apply the filter and display messages, click **Search**.

Table 10: Summary of Key Routing Information Output Fields

Field	Values	Additional Information
Static Route Addresses	The list of static route addresses.	
Protocol	Protocol from which the route was learned: **Static**, **Direct**, **Local**.	
Preference	The preference is the individual preference value for the route.	The route preference is used as one of the route selection criteria.

Copyright © 2017, Juniper Networks, Inc.

Table 10: Summary of Key Routing Information Output Fields *(continued)*

Field	Values	Additional Information
Next-Hop	Network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.	If a next hop is listed as **Discard**, all traffic with that destination address is discarded rather than routed. This value generally means that the route is a static route for which the **discard** attribute has been set. If a next hop is listed as **Reject**, all traffic with that destination address is rejected. This value generally means that the address is unreachable. For example, if the address is a configured interface address and the interface is unavailable, traffic bound for that address is rejected. If a next hop is listed as **Local**, the destination is an address on the host (either the loopback address or Ethernet management port 0 address, for example).
Age	How long the route has been active.	
State	Flags for this route.	There are many possible flags.

Related Documentation

- Configuring Static Routing on page 4

Verifying That Virtual Router Routing Instances Are Working

Purpose After creating a virtual router routing instance, verify that it has been set up properly.

Action

1. Use the **show route instance** command to list all the routing instances and their properties:

```
user@switch> show route instance
Instance             Type
        Primary RIB                                  Active/holddown/hidden
master               forwarding
        inet.0                                         4/0/1

__juniper_private1__ forwarding
        __juniper_private1__.inet.0                    1/0/3

__juniper_private2__ forwarding
        __juniper_private2__.inet.0                    0/0/1

__juniper_private3__ forwarding
        __juniper_private3__.inet.0                    1/0/2

__juniper_private4__ forwarding
        __juniper_private4__.inet.0                    4/0/2

__master.anon__      forwarding

r1                   virtual-router

r2                   virtual-router
```

Copyright © 2017, Juniper Networks, Inc.

2. Use the **show route forwarding-table** command to view the forwarding table information for each routing instance:

```
user@switch> show route forwarding-table
Routing table: r1---qfabric.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct  1628     1
0.0.0.0/32         perm     0                    dscd  1626     1
224.0.0.0/4        perm     0                    mdsc  1627     1
224.0.0.1/32       perm     0 224.0.0.1          mcst  1623     1
255.255.255.255/32 perm     0                    bcst  1624     1
```

Meaning The output displays the routing table information and confirms that the virtual router routing instances have been created and the links are up.

Related Documentation
- Configuring Virtual Router Routing Instances on page 8

Troubleshooting Virtual Routing Instances

- Direct Routes Not Leaked Between Routing Instances on page 48

Direct Routes Not Leaked Between Routing Instances

Problem **Description:** Direct routes are not exported (leaked) between virtual routing instances For example, consider the following scenario:

- Switch with two virtual routing instances:
  - Routing instance 1 connects to downstream device through interface xe-0/0/1.
  - Routing instance 2 connects to upstream device through interface xe-0/0/2.

If you enable route leaking between the routing instances (by using the **rib-group** statement, for example), the downstream device cannot connect to the upstream device because the switch connects to the upstream device over a direct route and these routes are not leaked between instances.

NOTE: You can see a route to the upstream device in the routing table of the downstream device, but this route is not functional.

Indirect routes *are* leaked between routing instances, so the downstream device can connect to any upstream devices that are connected to the switch over indirect routes.

Solution This is expected behavior.

Related Documentation
- Understanding Virtual Router Routing Instances on page 7
- Configuring Virtual Router Routing Instances on page 8

Copyright © 2017, Juniper Networks, Inc.

- rib-group on page 150

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

PART 2

Configuration Statements and Operational Commands

- Routing Options Configuration Statements on page 53
- Routing Options Operational Commands on page 175

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 2

Routing Options Configuration Statements

- active on page 56
- aggregate (Routing) on page 57
- as-path (Routing Options) on page 59
- autonomous-system on page 61
- backup-pe-group on page 63
- backups on page 64
- bandwidth on page 65
- bfd-liveness-detection (Routing Options Static Route) on page 66
- bgp-orf-cisco-mode on page 70
- bmp on page 72
- brief on page 74
- centralized on page 75
- community (Routing Options) on page 76
- confederation on page 78
- disable (Routing Options) on page 79
- description (Routing Instances) on page 79
- discard on page 80
- export (Routing Options) on page 81
- export-rib on page 82
- fate-sharing on page 85
- flow on page 86
- flow-map on page 87
- forwarding-cache (Flow Maps) on page 88
- forwarding-cache (Multicast) on page 88
- forwarding-options (chassis) on page 89
- forwarding-table on page 92

Copyright © 2017, Juniper Networks, Inc.

- generate on page 93
- graceful-restart (Enabling Globally) on page 95
- import on page 96
- import-policy on page 97
- import-rib on page 98
- indirect-next-hop on page 99
- install (Routing Options) on page 100
- instance-export on page 101
- instance-import on page 101
- instance-type on page 102
- interface (Multicast Static Routes) on page 103
- interface (Routing Instances) on page 104
- interface (Routing Options) on page 105
- interface-routes on page 106
- local-address (Routing Options) on page 107
- martians on page 108
- maximum-bandwidth on page 109
- maximum-paths on page 110
- maximum-prefixes on page 112
- med-igp-update-interval on page 113
- metric (Aggregate, Generated, or Static Route) on page 114
- multicast (Routing Options) on page 115
- no-qos-adjust on page 116
- num-65-127-prefix on page 117
- options (Routing Options) on page 118
- pim-to-igmp-proxy on page 119
- pim-to-mld-proxy on page 120
- policy (Aggregate and Generated Routes) on page 121
- policy (Flow Maps) on page 122
- policy-options on page 123
- policy-statement on page 124
- ppm on page 128
- ppm (Ethernet Switching) on page 129
- preference (Routing Options) on page 130
- prefix on page 131
- prefix-65-127-disable on page 132
- protocols on page 133

 Copyright © 2017, Juniper Networks, Inc.

- qualified-next-hop (Static Routes) on page 136
- readvertise on page 138
- redundant-sources on page 139
- resolution on page 140
- resolution-ribs on page 141
- resolve on page 142
- restart-duration (Routing Options) on page 143
- retain on page 144
- reverse-oif-mapping on page 145
- rpf-check-policy (Routing Options RPF) on page 146
- rib (General) on page 147
- rib (Route Resolution) on page 149
- rib-group (Routing Options) on page 150
- rib-groups on page 151
- route-distinguisher-id on page 153
- route-record on page 154
- router-id on page 155
- routing-instances on page 156
- routing-options on page 157
- scope on page 157
- scope-policy on page 158
- source (Source-Specific Multicast) on page 159
- source-routing on page 160
- ssm-groups on page 161
- ssm-map (Routing Options Multicast) on page 162
- static (Routes) on page 163
- subscriber-leave-timer on page 165
- tag (Routing Options) on page 166
- threshold (Multicast Forwarding Cache) on page 167
- timeout (Flow Maps) on page 168
- timeout (Multicast) on page 169
- traceoptions on page 170
- upstream-interface on page 173

Copyright © 2017, Juniper Networks, Inc.

active

Syntax (active | passive);

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Determine whether static, aggregate, or generated routes are removed from the routing and forwarding tables when they become inactive. Static routes are only removed from the routing table if the next hop becomes unreachable. This can occur if the local or neighbor interface goes down. Routes that have been configured to remain continually installed in the routing and forwarding tables are marked with **reject** next hops when they are inactive.

- **active**—Remove a route from the routing and forwarding tables when it becomes inactive.
- **passive**—Have a route remain continually installed in the routing and forwarding tables even when it becomes inactive.

Include the **active** statement when configuring an individual route in the **route** portion of the **static** statement to override a **passive** option specified in the **defaults** portion of the statement.

Default active

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Summarizing Static Routes Through Route Aggregation*
- *Example: Configuring a Conditional Default Route Policy*

Copyright © 2017, Juniper Networks, Inc.

aggregate (Routing)

Syntax

```
aggregate {
  defaults {
    ... aggregate-options ...
  }
  route destination-prefix {
    policy policy-name;
    ... aggregate-options ...
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name*],
[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name*],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name*],
[edit routing-options],
[edit routing-options rib *routing-table-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure aggregate routes.

Options

aggregate-options—Additional information about aggregate routes that is included with the route when it is installed in the routing table. Specify zero or more of the following options in ***aggregate-options***. Each option is explained separately.

- *active*(— Removes inactive routes from the forwarding table.
- ***passive***— Retains inactive routes in the forwarding table.
- as-path <*as-path*> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator *as-number ip-address*>;
- (brief | full);
- community [*community-ids*];
- discard;
- (metric | metric2 | metric3 | metric4) *value* <type *type*>;
- (preference | preference2 | color | color2) *preference* <type *type*>;
- tag *metric* type *number*;

Copyright © 2017, Juniper Networks, Inc.

defaults—Specify global aggregate route options. These options only set default attributes inherited by all newly created aggregate routes. These are treated as global defaults and apply to all the aggregate routes you configure in the **aggregate** statement. This part of the **aggregate** statement is optional.

route *destination-prefix*—Configure a nondefault aggregate route:

- **default**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.
- ***destination-prefix/prefix-length*—*destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.

The policy statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Summarizing Static Routes Through Route Aggregation*

Copyright © 2017, Juniper Networks, Inc.

as-path (Routing Options)

Syntax as-path <*as-path*> <aggregator *as-number ip-address*> <atomic-aggregate> <origin (egp | igp | incomplete)>;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Associate BGP autonomous system (AS) path information with a static, aggregate, or generated route.

In Junos OS Release 9.1 and later, the numeric range for the AS number is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.

In Junos OS Release 9.2 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <*16-bit high-order value in decimal*>.<*16-bit low-order value in decimal*>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format. You can specify a value in the range from 0.0 through 65535.65535 in AS-dot notation format.

Default No AS path information is associated with static routes.

Options **aggregator**—(Optional) Attach the BGP **aggregator** path attribute to the aggregate route. You must specify the last AS number that formed the aggregate route (encoded as two octets) for ***as-number***, followed by the IP address of the BGP system that formed the aggregate route for ***ip-address***.

Copyright © 2017, Juniper Networks, Inc.

as-path—(Optional) AS path to include with the route. It can include a combination of individual AS path numbers and AS sets. Enclose sets in brackets ([ ]). The first AS number in the path represents the AS immediately adjacent to the local AS. Each subsequent number represents an AS that is progressively farther from the local AS, heading toward the origin of the path. You cannot specify a regular expression for ***as-path***. You must use a complete, valid AS path.

atomic-aggregate—(Optional) Attach the BGP **atomic-aggregate** path attribute to the aggregate route. This path attribute indicates that the local system selected a less specific route instead of a more specific route.

origin egp—(Optional) BGP origin attribute that indicates that the path information originated in another AS.

origin igp—(Optional) BGP origin attribute that indicates that the path information originated within the local AS.

origin incomplete—(Optional) BGP origin attribute that indicates that the path information was learned by some other means.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Summarizing Static Routes Through Route Aggregation*
- *Using 4-Byte Autonomous System Numbers in BGP Networks Technology Overview*

Copyright © 2017, Juniper Networks, Inc.

autonomous-system

Syntax

```
autonomous-system autonomous-system <asdot-notation> <loops number> {
  independent-domain <no-attrset>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
asdot-notation option introduced in Junos OS Release 9.3.
asdot-notation option introduced in Junos OS Release 9.3 for EX Series switches.
no-attrset option introduced in Junos OS Release 10.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Specify the routing device's AS number.

An autonomous system (AS) is a set of routing devices that are under a single technical administration and that generally use a single interior gateway protocol (IGP) and metrics to propagate routing information within the set of routing devices. An AS appears to other ASs to have a single, coherent interior routing plan and presents a consistent picture of what destinations are reachable through it. ASs are identified by a number that is assigned by the Network Information Center (NIC) in the United States (http://www.isi.edu).

If you are using BGP on the routing device, you must configure an AS number.

The AS path attribute is modified when a route is advertised to an EBGP peer. Each time a route is advertised to an EBGP peer, the local routing device prepends its AS number to the existing path attribute, and a value of 1 is added to the AS number.

In Junos OS Release 9.1 and later, the numeric range is extended to provide BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*. RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. RFC 4893 also introduces a reserved, well-known, 2-byte AS number, AS 23456. This reserved AS number is called AS_TRANS in RFC 4893. All releases of Junos OS support 2-byte AS numbers.

In Junos OS Release 9.3 and later, you can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: <*16-bit high-order value in decimal*>.<*16-bit low-order value in decimal*>. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.

Options

autonomous-system—AS number. Use a number assigned to you by the NIC.

Copyright © 2017, Juniper Networks, Inc.

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format for 4-byte AS numbers

In this example, the 4-byte AS number 65,546 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 65546;
}
```

Range: 0.0 through 65535.65535 in AS-dot notation format for 4-byte numbers

In this example, 1.10 is the AS-dot notation format for 65,546:

```
[edit]
routing-options {
  autonomous-system 1.10;
}
```

Range: 1 through 65,535 in plain-number format for 2-byte AS numbers (this is a subset of the 4-byte range)

In this example, the 2-byte AS number 60,000 is represented in plain-number format:

```
[edit]
routing-options {
  autonomous-system 60000;
}
```

asdot-notation—(Optional) Display the configured 4-byte autonomous system number in the AS-dot notation format.

Default: Even if a 4-byte AS number is configured in the AS-dot notation format, the default is to display the AS number in the plain-number format.

loops ***number***—(Optional) Specify the number of times detection of the AS number in the AS_PATH attribute causes the route to be discarded or hidden. For example, if you configure **loops 1**, the route is hidden if the AS number is detected in the path one or more times. This is the default behavior. If you configure **loops 2**, the route is hidden if the AS number is detected in the path two or more times.

Range: 1 through 10

Default: 1

NOTE: When you specify the same AS number in more than one routing instance on the local routing device, you must configure the same number of loops for the AS number in each instance. For example, if you configure a value of 3 for the **loops** statement in a VRF routing instance that uses the same AS number as that of the master instance, you must also configure a value of 3 loops for the AS number in the master instance.

Use the **independent-domain** option if the **loops** statement must be enabled only on a subset of routing instances.

The remaining statement is explained separately.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Examples: Configuring External BGP Peering*
- *Examples: Configuring Internal BGP Peering*

backup-pe-group

Syntax

```
backup-pe-group group-name {
  backups [ addresses ];
  local-address address;
}
```

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information
Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description
Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.

Options
group-name—Name of the group for PE backups.

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring Ingress PE Redundancy*

Copyright © 2017, Juniper Networks, Inc.

backups

Syntax backups [*addresses*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast backup-pe-group *group-name*],
[edit logical-systems *logical-system-name* routing-options multicast backup-pe-group *group-name*],
[edit routing-instances *routing-instance-name* routing-options multicast backup-pe-group *group-name*],
[edit routing-options multicast backup-pe-group *group-name*]

Release Information Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the address of backup PEs for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.

Options ***addresses***—Addresses of other PEs in the backup group.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring Ingress PE Redundancy*

Copyright © 2017, Juniper Networks, Inc.

bandwidth

Syntax bandwidth (*bps* | adaptive);

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map],
[edit logical-systems *logical-system-name* routing-options multicast flow-map],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map],
[edit routing-options multicast flow-map]

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure the bandwidth property for multicast flow maps.

Options **adaptive**—Specify that the bandwidth is measured for the flows that are matched by the flow map.

bps—Bandwidth, in bits per second, for the flow map.

Range: 0 through any amount of bandwidth

Default: 2 Mbps

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring a Multicast Flow Map*

Copyright © 2017, Juniper Networks, Inc.

bfd-liveness-detection (Routing Options Static Route)

Syntax

```
bfd-liveness-detection {
    description Site- xxx;
    authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
    }
    detection-time {
        threshold milliseconds;
    }
    holddown-interval milliseconds;
    local-address ip-address;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    minimum-receive-ttl number;
    multiplier number;
    neighbor address;
    no-adaptation;
    transmit-interval {
        minimum-interval milliseconds;
        threshold milliseconds;
    }
    version (1 | automatic);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit logical-systems *logical-system-name* routing-options static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix*],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit routing-instances *routing-instance-name* routing-options static route *destination-prefix*],
[edit routing-instances *routing-instance-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit routing-options rib *routing-table-name* static route *destination-prefix*],
[edit routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],

Copyright © 2017, Juniper Networks, Inc.

[edit routing-options static route *destination-prefix*],
[edit routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)]

Release Information Statement introduced before Junos OS Release 7.4.
detection-time threshold and **transmit-interval threshold** options introduced in Junos OS Release 8.2.
local-address statement introduced in Junos OS Release 8.2.
minimum-receive-ttl statement introduced in Junos OS Release 8.2.
Support for logical routers introduced in Junos OS Release 8.3.
holddown-interval statement introduced in Junos OS Release 8.5.
no-adaptation statement introduced in Junos OS Release 9.0.
Support for IPv6 static routes introduced in Junos OS Release 9.1.
authentication algorithm, **authentication key-chain**, and **authentication loose-check** statements introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure bidirectional failure detection timers and authentication criteria for static routes.

Copyright © 2017, Juniper Networks, Inc.

Options

authentication algorithm ***algorithm-name*** —Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

authentication key-chain ***key-chain-name***—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold ***milliseconds***—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

holddown-interval ***milliseconds***—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent. If the BFD session goes down and then comes back up during the hold-down interval, the timer is restarted.

Range: 0 through 255,000

Default: 0

local-address ***ip-address***—Enable a multihop BFD session and configure the source address for the BFD session.

minimum-interval ***milliseconds***—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

Range: 1 through 255,000

minimum-receive-interval ***milliseconds***—Configure the minimum interval after which the routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement at the **[edit routing-options static route** ***destination-prefix*** **bfd-liveness-detection]** hierarchy level.

Range: 1 through 255,000

minimum-receive-ttl ***number***—Configure the time to live (TTL) for the multihop BFD session.

Range: 1 through 255

Default: 255

multiplier ***number***—Configure number of hello packets not received by the neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

Copyright © 2017, Juniper Networks, Inc.

neighbor *address*—Configure a next-hop address for the BFD session for a next hop specified as an interface name.

no-adaptation—Specify for BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds*—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295

transmit-interval minimum-interval *milliseconds*—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement at the **[edit routing-options static route *destination-prefix* bfd-liveness-detection]** hierarchy level.

Range: 1 through 255,000

version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

Default: automatic

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring BFD for Static Routes for Faster Network Failure Detection on page 16
- Example: Configuring BFD Authentication for Securing Static Routes on page 30

Copyright © 2017, Juniper Networks, Inc.

bgp-orf-cisco-mode

Syntax bgp-orf-cisco-mode;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp outbound-route-filter],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* outbound-route-filter],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address* outbound-route-filter],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp outbound-route-filter],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* outbound-route-filter,
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* outbound-route-filter],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options outbound-route-filter],
[edit logical-systems *logical-system-name* routing-options outbound-route-filter],
[edit protocols bgp outbound-route-filter],
[edit protocols bgp group *group-name* outbound-route-filter],
[edit protocols bgp group *group-name* neighbor *address* outbound-route-filter],
[edit routing-instances *routing-instance-name* protocols bgp outbound-route-filter],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* outbound-route-filter],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* outbound-route-filter],
[edit routing-instances *routing-instance-name* routing-options outbound-route-filter],
[edit routing-options outbound-route-filter]

Release Information
Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 9.2 for EX Series switches.
Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.2.
Support for the BGP group and neighbor hierarchy levels introduced in Junos OS Release 9.3 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Enable interoperability with routing devices that use the vendor-specific outbound route filter compatibility code of 130 and code type of 128.

NOTE: To enable interoperability for all BGP peers configured on the routing device, include the statement at the **[edit routing-options outbound-route-filter]** hierarchy level.

Default Disabled

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring BGP Prefix-Based Outbound Route Filtering*

Copyright © 2017, Juniper Networks, Inc.

bmp

Syntax

```
bmp {
  authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
  authentication-key key;
  authentication-key-chain authentication-key-chain;
  connection-mode (active | passive);
  hold-down {
    seconds;
    flaps flaps;
    period seconds;
  }
  initiation-message text;
  local-address address;
  local-port port;
  monitor (disable | enable);
  priority (high | low | medium);
  route-monitoring {
    none;
    post-policy {
      exclude-non-eligible;
    }
    pre-policy {
      exclude-non-feasible;
    }
  }
  station station-name {
    authentication-algorithm (aes-128-cmac-96 | hmac-sha-1-96 | md5);
    authentication-key key;
    authentication-key-chain authentication-key-chain;
    connection-mode (active | passive);
    hold-down {
      seconds;
      flaps flaps;
      period seconds;
    }
    initiation-message text;
    local-address address;
    local-port port;
    monitor (disable | enable);
    priority (high | low | medium);
    route-monitoring {
      none;
      post-policy {
        exclude-non-eligible;
      }
      pre-policy {
        exclude-non-feasible;
      }
    }
    station-address (ip-address | name);
    station-port port-number;
    statistics-timeout seconds;
    traceoptions {
```


Copyright © 2017, Juniper Networks, Inc.

```
        file filename <files number> <size size> <world-readable | no-world-readable>;
        flag flag <flag-modifier>;
      }
    }
    station-address (ip-address | name);
    station-port port-number;
    statistics-timeout seconds;
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier>;
    }
}
```

Hierarchy Level

[edit routing-options]
[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp group *group-name*],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address*],
[edit protocols bgp],
[edit protocols bgp group *group-name*],
[edit protocols bgp group *group-name* neighbor *address*],

NOTE: 1. Complete BMP configuration, as mentioned in the syntax, can be done under the first two hierarchy levels only

2. Under other hierarchy levels, only the following configurations are supported:

- Either we can inherit or not inherit the configuration data
- Enable/disable monitoring
- Control route monitoring settings

Release Information

Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Support for BMP version 3 introduced in Junos OS Release 13.3.

Description

Configure the BGP Monitoring Protocol (BMP), which enables the routing device to collect data from the BGP Adjacency-RIB-In routing tables and periodically send that data to a monitoring station.

Options

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring the BGP Monitoring Protocol*

Copyright © 2017, Juniper Networks, Inc.

brief

Syntax (brief | full);

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (aggregate | generate) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit routing-options (aggregate | generate) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure all AS numbers from all contributing paths to be included in the aggregate or generated route's path.

- **brief**—Include only the longest common leading sequences from the contributing AS paths. If this results in AS numbers being omitted from the aggregate route, the BGP **ATOMIC_ATTRIBUTE** path attribute is included with the aggregate route.
- **full**—Include all AS numbers from all contributing paths in the aggregate or generated route's path. Include this option when configuring an individual route in the **route** portion of the **generate** statement to override a **retain** option specified in the **defaults** portion of the statement.

Default full

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Summarizing Static Routes Through Route Aggregation*
- *Example: Configuring a Conditional Default Route Policy*
- *Understanding Conditionally Generated Routes*
- aggregate on page 57
- generate on page 93

Copyright © 2017, Juniper Networks, Inc.

centralized

Syntax centralized;

Hierarchy Level [edit protocols lacp ppm]

Release Information Statement introduced in Junos OS Release 9.4 for MX Series routers.
Statement introduced in Junos OS Release 10.2 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Disable distributed periodic packet management (PPM) processing for Link Aggregation Control Protocol (LACP) packets and run all PPM processing for LACP packets on the Routing Engine.

This statement disables distributed PPM processing for only LACP packets. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the **no-delegate-processing** statement in the **[edit routing-options** ppm**]** hierarchy.

BEST PRACTICE: We generally recommend that you disable distributed PPM only if Juniper Networks Customer Service advises you to do so. You should disable distributed PPM only if you have a compelling reason to disable it.

Default Distributed PPM processing is enabled for all packets that use PPM.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)*
- *Configuring Aggregated Ethernet LACP (CLI Procedure)*
- Configuring Distributed Periodic Packet Management on page 9
- *Configuring Link Aggregation*

Copyright © 2017, Juniper Networks, Inc.

community (Routing Options)

Syntax community ([*community-ids*] | no-advertise | no-export | no-export-subconfed | none | llgr-stale | no-llgr);

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
llgr-stale and **no-llgr** options added in Junos OS Release 15.1.

Description Associate BGP community information with a static, aggregate, or generated route.

Default No BGP community information is associated with static routes.

Options ***community-ids***—One or more community identifiers. The ***community-ids*** format varies according to the type of attribute that you use.

The BGP community attribute format is ***as-number:community-value***:

- ***as-number***—AS number of the community member. It can be a value from 1 through 65,535. The AS number can be a decimal or hexadecimal value.
- ***community-value***—Identifier of the community member. It can be a number from 0 through 65,535.

For more information about BGP community attributes, see the "Configuring the Extended Communities Attribute" section in the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

For specifying the BGP community attribute only, you also can specify ***community-ids*** as one of the following well-known community names defined in RFC 1997:

- **no-advertise**—Routes containing this community name are not advertised to other BGP peers.

Copyright © 2017, Juniper Networks, Inc.

- **no-export**—Routes containing this community name are not advertised outside a BGP confederation boundary.
- **no-export-subconfed**—Routes containing this community are advertised to IBGP peers with the same AS number, but not to members of other confederations.
- **llgr-stale**—Adds a community to a long-lived stale route when it is readvertised.
- **no-llgr**—Marks routes which a BGP speaker does not want to be retained by LLGR. The Notification message feature does not have any associated configuration parameters.

NOTE: Extended community attributes are not supported at the **[edit routing-options]** hierarchy level. You must configure extended communities at the **[edit policy-options]** hierarchy level. For information about configuring extended communities, see the *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Summarizing Static Routes Through Route Aggregation*
- aggregate on page 57
- generate on page 93
- *static*

Copyright © 2017, Juniper Networks, Inc.

confederation

Syntax confederation *confederation-autonomous-system* members [*autonomous-systems*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Specify the routing device's confederation AS number.

If you administer multiple ASs that contain a very large number of BGP systems, you can group them into one or more *confederations*. Each confederation is identified by its own AS number, which is called a *confederation AS number*. To external ASs, a confederation appears to be a single AS. Thus, the internal topology of the ASs making up the confederation is hidden.

The BGP path attributes **NEXT_HOP**, **LOCAL_PREF**, and **MULTI_EXIT_DISC**, which normally are restricted to a single AS, are allowed to be propagated throughout the ASs that are members of the same confederation.

Because each confederation is treated as if it were a single AS, you can apply the same routing policy to all the ASs that make up the confederation.

Grouping ASs into confederations reduces the number of BGP connections required to interconnect ASs.

If you are using BGP, you can enable the local routing device to participate as a member of an AS confederation. To do this, include the **confederation** statement.

Specify the AS confederation identifier, along with the peer AS numbers that are members of the confederation.

Note that peer adjacencies do not form if two BGP neighbors disagree about whether an adjacency falls within a particular confederation.

Options ***autonomous-systems***—AS numbers of the confederation members.

Range: 1 through 65,535

confederation-autonomous-system—Confederation AS number. Use one of the numbers assigned to you by the NIC.

Range: 1 through 65,535

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation
- *Example: Configuring BGP Confederations*
- *Understanding BGP Confederations*

disable (Routing Options)

Syntax disable;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options graceful-restart],
[edit logical-systems *logical-system-name* routing-options graceful-restart],
[edit routing-instances *routing-instance-name* routing-options graceful-restart],
[edit routing-options graceful-restart]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Disable graceful restart.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Junos OS High Availability Library for Routing Devices*

description (Routing Instances)

Syntax description *text*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],
[edit routing-instances *routing-instance-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.1 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Provide a text description for the routing instance. If the text includes one or more spaces, enclose it in quotation marks (" "). Any descriptive text you include is displayed in the output of the **show route instance detail** command and has no effect on the operation of the routing instance.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

discard

Syntax discard;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (aggregate | generate) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit routing-options (aggregate | generate) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

To propagate static routes into the routing protocols, include the **discard** statement when you define the route, along with a routing policy.

NOTE: In other vendors' software, a common way to propagate static routes into routing protocols is to configure the routes so that the next-hop routing device is the loopback address (commonly, 127.0.0.1). However, configuring static routes in this way (by including a statement such as **route *address/mask-length* next-hop 127.0.0.1**) does not propagate the static routes, because the forwarding table ignores static routes whose next-hop routing device is the loopback address.

Default When an aggregate route becomes active, it is installed in the routing table with a reject next hop, which means that ICMP unreachable messages are sent.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Summarizing Static Routes Through Route Aggregation*

Copyright © 2017, Juniper Networks, Inc.

- aggregate on page 57
- generate on page 93

export (Routing Options)

Syntax export [*policy-name*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options forwarding-table],
[edit logical-systems *logical-system-name* routing-options forwarding-table],
[edit routing-instances *routing-instance-name* routing-options forwarding-table],
[edit routing-options forwarding-table]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Apply one or more policies to routes being exported from the routing table into the forwarding table.

In the **export** statement, list the name of the routing policy to be evaluated when routes are being exported from the routing table into the forwarding table. Only active routes are exported from the routing table.

You can reference the same routing policy one or more times in the same or a different **export** statement.

You can apply export policies to routes being exported from the routing table into the forwarding table for the following features:

- Per-packet load balancing
- Class of service (CoS)

Options ***policy-name***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Load Balancing BGP Traffic*

Copyright © 2017, Juniper Networks, Inc.

export-rib

Syntax export-rib *routing-table-name*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib-groups *group-name*],
[edit logical-systems *logical-system-name* routing-options rib-groups *group-name*],
[edit routing-instances *routing-instance-name* routing-options rib-groups *group-name*],
[edit routing-options rib-groups *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Specify the name of the routing table from which Junos OS should export routing information. For any individual RIB group, only one table can be specified in the **export-rib** statement.

The **export-rib** statement specifies the source table from which routing information is advertised.

One common use of the **export-rib** statement is interdomain routing. The export RIB is the table used when BGP extracts routes to advertise to peers. In multicast interdomain routing, for example, the export RIB is likely to be inet.2.

Another use of **export-rib** is dynamic route leaking between the global routing table (inet.0) and a VRF routing table (*instance*.inet.0). For example, you can use a RIB group to copy routes learned in the VRF into the global routing table, inet.0, or copy routes learned in inet.0 into a VRF. You define the use of this RIB group in the VRF's BGP configuration. In a routing policy you can do dynamic filtering of routes. For instance, you can use an import policy to only copy routes with certain communities into the global routing table.

For example:

```
rib-groups {
    rib-interface-routes-v4 {
        import-rib [ inet.0 VRF.inet.0 ];
    }
    rib-import-VRF-routes-to-inet0-v4 {
        export-rib VRF.inet.0;
        import-rib [ VRF.inet.0 inet.0 ];
        import-policy rib-import-VRF-routes-to-inet0-v4;
    }
    rib-import-inet0-routes-to-VRF-v4 {
        export-rib inet.0;
        import-rib [ inet.0 VRF.inet.0 ];
        import-policy rib-import-inet0-routes-to-VRF-v4;
    }
}
routing-options {
    interface-routes {
```


Copyright © 2017, Juniper Networks, Inc.

```
      rib-group {
        inet rib-interface-routes-v4;
      }
    }
  }
  protocols {
    bgp {
      group iBGP-peers {
        type internal;
        family inet {
          unicast {
            rib-group rib-import-inet0-routes-to-VRF-v4;
          }
        }
      }
    }
  }
  routing-instances {
    VRF {
      routing-options {
        interface-routes {
          rib-group {
            inet rib-interface-routes-v4;
          }
        }
      }
      protocols {
        bgp {
          group peersin-VRF {
            family inet {
              unicast {
                rib-group rib-import-VRF-routes-to-inet0-v4;
              }
            }
          }
        }
      }
    }
  }
```

Options *routing-table-name*—Routing table group name.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- *Example: Exporting Specific Routes from One Routing Table Into Another Routing Table*
- *Example: Configuring a PIM RPF Routing Table*
- *Example: Configuring DVMRP to Announce Unicast Routes*
- *Example: Configuring a Dedicated PIM RPF Routing Table*
- *Example: Configuring Any-Source Multicast for Draft-Rosen VPNs*
- import-rib on page 98
- *passive*

Copyright © 2017, Juniper Networks, Inc.

fate-sharing

Syntax

```
fate-sharing {
  group group-name {
    cost value;
    from address <to address>;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit routing-options],
[edit routing-instances *routing-instance-name* routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Specify a backup path in case the primary path becomes unusable.

You specify one or more objects with common characteristics within a group. All objects are treated as /32 host addresses. The objects can be a LAN interface, a router ID, or a point-to-point link. Sequence is insignificant.

Changing the fate-sharing database does not affect existing established LSPs until the next CSPF reoptimization. The fate-sharing database does affect fast-reroute detour path computations.

Options

cost *value*—Cost assigned to the group.

Range: 1 through 65,535

Default: 1

from *address*—Address of the router or address of the LAN/NBMA interface. For example, an Ethernet network with four hosts in the same fate-sharing group would require you to list all four of the separate **from** addresses in the group.

group *group-name*—Each fate-sharing group must have a name, which can have a maximum of 32 characters, including letters, numbers, periods (.), and hyphens (-). You can define up to 512 groups.

to *address*—(Optional) Address of egress router. For point-to-point link objects, you must specify both a **from** and a **to** address.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring the Ingress Router for MPLS-Signaled LSPs*
- *MPLS Applications Feature Guide*

Copyright © 2017, Juniper Networks, Inc.

flow

Syntax

```
flow {
  route name {
    match {
      match-conditions;
    }
    term-order (legacy | standard);
    then {
      actions;
    }
  }
  firewall-install-disable;
  term-order (legacy | standard);
  validation {
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
    }
  }
}
```

Hierarchy Level

[edit routing-options],
[edit routing-instances *routing-instance-name* routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
term-order statement introduced in Junos OS Release 10.0
Statement introduced in Junos OS Release 11.3 for the QFX Series.
firewall-install-disable statement introduced in Junos OS Releases 12.1X48 and 12.3 for PTX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure a flow route.

Default

legacy

Options

actions—An action to take if conditions match.

firewall-install-disable—(PTX Series routers only) Disable installing flow-specification firewall filters in the firewall process (dfwd).

Default: If you omit the **firewall-install-disable** statement, the default behavior is **firewall-install-disable** mode.

match-conditions—Match packets to these conditions.

route *name*—Name of the flow route.

standard—Specify to use version 7 or later of the flow-specification algorithm.

term-order (legacy | standard)—Specify the version of the flow-specification algorithm.

Copyright © 2017, Juniper Networks, Inc.

- **legacy**—Use version 6 of the flow-specification algorithm.
- **standard**—Use version 7 of the flow-specification algorithm.

then—Actions to take on matching packets.

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Enabling BGP to Carry Flow-Specification Routes*
- *Understanding BGP Flow Routes for Traffic Filtering*

flow-map

Syntax

```
flow-map flow-map-name {
  bandwidth (bps | adaptive);
  forwarding-cache {
    timeout (never non-discard-entry-only | minutes);
  }
  policy [ policy-names ];
  redundant-sources [ addresses ];
}
```

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information
Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description
Configure multicast flow maps.

Options
flow-map-name—Name of the flow-map.

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring a Multicast Flow Map*

Copyright © 2017, Juniper Networks, Inc.

forwarding-cache (Flow Maps)

Syntax

```
forwarding-cache {
  timeout (minutes | never non-discard-entry-only );
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-options multicast flow-map *flow-map-name*]

Release Information

Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure multicast forwarding cache properties for the flow map.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring a Multicast Flow Map*

forwarding-cache (Multicast)

Syntax

```
forwarding-cache {
  threshold suppress value <reuse value>;
  timeout minutes;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure multicast forwarding cache properties. These properties include threshold suppression and reuse limits and timeout values.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring the Multicast Forwarding Cache*

Copyright © 2017, Juniper Networks, Inc.

forwarding-options (chassis)

Syntax

```
forwarding options profile-name {
  num-65-127-prefix number;
}
```

Syntax

```
forwarding-options lpm-profile {
  prefix-65-127-disable;
  unicast-in-lpm;
}
```

Syntax

```
forwarding-options custom-profile {
  l2-entries | l3-entries | lpm-entries {
    num-banks number;
  }
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 13.2.
custom-profile option introduced in Junos OS Release 15.1x53-D30 for QFX5200 Series switches only.

Description Configure a unified forwarding table profile to allocate the amount of memory available for the following:

- MAC addresses.
- Layer 3 host entries.
- Longest prefix match table entries.

This feature enables you to select a profile that optimizes the amount of memory available for various types of forwarding-table entries based on the needs of your network. For example, for a switch that handles a great deal of Layer 2 traffic, such as a virtualized network with many servers and virtualized machines, you would choose the **l2-profile-one**, which allocates the highest amount of memory to MAC addresses.

You configure the memory allocation for LPM table entries differently, depending on whether you using Junos OS Release 13.2X51-D10 or Junos OS Release 13.2X51-D15 and later. For more information about configuring memory allocation for LPM table entries, see "Configuring the Unified Forwarding Table" on page 39.

The **num-65-127-prefix** *number* statement is not supported on the **custom-profile** and the **lpm-profile**. The **prefix-65-127-disable** and **unicast-in-lpm** statements are supported only on the **lpm-profile**.

When you commit a configuration with a forwarding table profile change, in most cases the Packet Forwarding Engine restarts automatically to apply the new parameters, which brings the data interfaces down and then up again.

However, starting with Junos OS Releases 14.1X53-D40, 15.1R5, and 16.1R3, for a Virtual Chassis or Virtual Chassis Fabric (VCF) comprised of EX4600 or QFX5100 switches, the

Copyright © 2017, Juniper Networks, Inc.

Packet Forwarding Engine in member switches does not automatically restart upon configuring and committing a unified forwarding table profile change. This behavior avoids having Virtual Chassis or VCF instability and a prolonged convergence period if a profile change is propagated to member switches and multiple Packet Forwarding Engines all restart at the same time. In this environment, instead of automatically restarting when you initially commit a profile configuration change, the message **Reboot required for configuration to take effect** is displayed at the master switch CLI prompt, notifying you that the profile change does not take effect until the next time you restart the Virtual Chassis or VCF. The profile configuration change is propagated to member switches that support this feature, and a reminder that a reboot is required to apply this pending configuration change appears in the system log of the master switch and applicable member switches. You then enable the profile change subsequently during a planned downtime period using the **request system reboot** command, which quickly establishes a stable Virtual Chassis or VCF with the new configuration.

NOTE: You should plan to make unified forwarding table profile changes only when you are ready to perform a Virtual Chassis or VCF system reboot *immediately* after committing the configuration update. Otherwise, in the intervening period between committing the configuration change and rebooting the Virtual Chassis or VCF, the system can become inconsistent if a member experiences a problem and restarts. In that case, the new configuration takes effect on the member that was restarted, while the change is not yet activated on all the other members.

Copyright © 2017, Juniper Networks, Inc.

Options **profile-name**—name of the profile to use for memory allocation in the unified forwarding table. Table 11 on page 91 lists the profiles you can choose that have set values and the associated values for each type of entry.

On QFX5200 Series switches only, you can also select **custom-profile**. This profile enables you to allocate from one to four banks of shared hash memory to a specific type of forwarding-table entry. Each shared hash memory bank can store a maximum of the equivalent of 32,000 IPv4 unicast addresses.

Table 11: Unified Forwarding Table Profiles

Profile Name	MAC Table	Host Table (unicast and multicast addresses)					
	MAC Addresses	IPv4 unicast	IPv6 unicast	IPv4 (*, G)	IPv4 (S, G)	IPv6 (*, G)	IPv6 (S, G)
l2-profile-one	288K	16K	8K	8K	8K	4K	4K
l2-profile-two	224K	80K	40K	40K	40K	20K	20K
l2-profile-three (default)	160K	144K	72K	72K	72K	36K	36K
l3-profile	96K	208K	104K	104K	104K	52K	52K
lpm-profile*	32K	16K	8K	8K	8K	4K	4K

* This profile supports only IPv4 in Junos OS Release13.2X51-D10. Starting in Junos OS Relase13.2X51-D15, the **lpm-profile** supports IPv4 and IPv6 entries.

NOTE: If the host stores the maximum number of entries for any given type, the entire table is full and is unable to accommodate *any* entries of any other type. For information about valid combinations of table entries see "Understanding the Unified Forwarding Table" on page 36.

l2-entries | l3-entries | lpm-entries—(**custom-profile only**) Select a type of forwarding-table entry—Layer 2, Layer 3, or LPM—to allocate a specific number of shared memory banks. You configure the amount of memory to allocate for each type of entry separately.

num-banks *number*—(**custom-profile only**) Specify the number of shared memory banks to allocate for a specific type of forwarding-table entry. Each shared memory bank stores the equivalent of 32,000 IPv4 unicast addresses.

Range: 0 through 4.

NOTE: There are four shared memory banks, which can be allocated flexibly among the three types of forwarding-table entries. To allocate no shared memory for a particular entry type, specify the number 0. When

Copyright © 2017, Juniper Networks, Inc.

you commit the configuration, the system issues a commit check to ensure that you have not configured more than four memory banks. You do not have to configure all four shared memory banks. By default, each entry type is allocated the equivalent of 32,000 IPv4 unicast addresses in shared memory.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- *Understanding the Unified Forwarding Table*
- *Example: Configuring a Unified Forwarding Table Custom Profile on QFX Series Switches*

forwarding-table

Syntax

```
forwarding-table {
  export [ policy--names ];
  (indirect-next-hop | no-indirect-next-hop);
}
```

Hierarchy Level
[edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Configure information about the routing device's forwarding table.

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Per-Packet Load Balancing on page 4

Copyright © 2017, Juniper Networks, Inc.

generate

Syntax

```
generate {
  defaults {
    generate-options;
  }
  route destination-prefix {
    policy policy-name;
    generate-options;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name*],
[edit routing-options],
[edit routing-options rib *routing-table-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure generated routes, which are used as routes of last resort.

Options

defaults—(Optional) Specify global generated route options. These options only set default attributes inherited by all newly created generated routes. These are treated as global defaults and apply to all the generated routes you configure in the **generate** statement.

generate-options—Additional information about generated routes, which is included with the route when it is installed in the routing table. Specify zero or more of the following options in ***generate-options***. Each option is explained separately.

- **(**active **| passive);**
- as-path **<***as-path***> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator** *as-number in-address***>;**
- **(**brief **| full);**
- community **[** *community-ids* **];**
- **discard;**
- **(**metric **| metric2 | metric3 | metric4)** *value* **<type** *type***>;**
- **(**preference **| preference2 | color | color2)** *preference* **<type** *type***>;**
- tag ***metric*** **type** ***number*****;**

route ***destination-prefix***—Configure a non-default generated route:

Copyright © 2017, Juniper Networks, Inc.

- **default**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.
- ***destination-prefix/prefix-length***—/***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.

The policy statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

graceful-restart (Enabling Globally)

Syntax

```
graceful-restart {
  disable;
  helper-disable;
  maximum-helper-recovery-time seconds;
  maximum-helper-restart-time seconds;
  notify-duration seconds;
  recovery-time seconds;
  restart-duration seconds;
  stale-routes-time seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit routing-options],
[edit routing-instances *routing-instance-name* routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure graceful restart globally to enable the feature. You cannot enable graceful restart for specific protocols unless graceful restart is also enabled globally. You can, optionally, modify the global settings at the individual protocol level.

NOTE:

- For VPNs, the **graceful-restart** statement allows a router whose VPN control plane is undergoing a restart to continue to forward traffic while recovering its state from neighboring routers.
- For BGP, if you configure graceful restart after a BGP session has been established, the BGP session restarts and the peers negotiate graceful restart capabilities.
- LDP sessions flap when **graceful-restart** configurations change.

Default

Graceful restart is disabled by default.

Options

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Enabling Graceful Restart*
- *Configuring Routing Protocols Graceful Restart*

Copyright © 2017, Juniper Networks, Inc.

- *Configuring Graceful Restart for MPLS-Related Protocols*
- *Configuring VPN Graceful Restart*
- *Configuring Logical System Graceful Restart*
- *Configuring Graceful Restart for QFabric Systems*

import

Syntax import [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options resolution rib],
[edit logical-systems *logical-system-name* routing-options resolution rib],
[edit routing-instances *routing-instance-name* routing-options resolution rib],
[edit routing-options resolution rib]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Specify one or more import policies to use for route resolution.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring Route Resolution on PE Routers*

Copyright © 2017, Juniper Networks, Inc.

import-policy

Syntax import-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib-groups *group-name*],
[edit logical-systems *logical-system-name* routing-options rib-groups *group-name*],
[edit routing-instances *routing-instance-name* routing-options rib-groups *group-name*],
[edit routing-options rib-groups *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Apply one or more policies to routes imported into the routing table group. The **import-policy** statement complements the **import-rib** statement and cannot be used unless you first specify the routing tables to which routes are being imported.

NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Exporting Specific Routes from One Routing Table Into Another Routing Table*
- export-rib on page 82
- *passive*

Copyright © 2017, Juniper Networks, Inc.

import-rib

Syntax import-rib [*routing-table-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib-groups *group-name*],
[edit logical-systems *logical-system-name* routing-options rib-groups *group-name*],
[edit routing-instances *routing-instance-name* routing-options rib-groups *group-name*],
[edit routing-options rib-groups *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Specify the name of the routing table into which Junos OS should import routing information. The first routing table name you enter is the primary routing table. Any additional names you enter identify secondary routing tables. When a protocol imports routes, it imports them into the primary and any secondary routing tables. If the primary route is deleted, the secondary route also is deleted. For IPv4 import routing tables, the primary routing table must be **inet.0** or **routing-instance-name.inet.0**. For IPv6 import routing tables, the primary routing table must be **inet6.0**.

In Junos OS Release 9.5 and later, you can configure an IPv4 import routing table that includes both IPv4 and IPv6 routing tables. Including both types of routing tables permits you, for example, to populate an IPv6 routing table with IPv6 addresses that are compatible with IPv4. In releases prior to Junos OS Release 9.5, you could configure an import routing table with only either IPv4 or IPv6 routing tables.

NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

Options ***routing-table-names***—Name of one or more routing tables.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Exporting Specific Routes from One Routing Table Into Another Routing Table*
- export-rib on page 82
- *passive*

Copyright © 2017, Juniper Networks, Inc.

indirect-next-hop

Syntax (indirect-next-hop | no-indirect-next-hop);

Hierarchy Level [edit logical-systems *logical-system-name* routing-options forwarding-table],
[edit routing-options forwarding-table]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 14.1X53-D30 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Enable indirectly connected next hops for route convergence. This statement is implemented on the Packet Forward Engine to speed up forwarding information base (FIB) updates. Configuring this statement significantly speeds convergence times. The only downside of configuring this statement is that some additional FIB memory overhead is required. Unless routes have an extremely high number of next hops, this increased memory usage should not be noticeable.

NOTE:

- When virtual private LAN service (VPLS) is configured on the routing device, the **indirect-next-hop** statement is configurable at the **[edit routing-options forwarding-table]** hierarchy level. However, this configuration is not applicable to indirect nexthops specific to VPLS routing instances.
- By default, the Junos Trio Modular Port Concentrator (MPC) chipset on MX Series routers is enabled with indirectly connected next hops, and this cannot be disabled using the **no-indirect-next-hop** statement.
- By default, indirectly connected next hops are enabled on PTX Series routers.

Default Disabled.

Options **indirect-next-hop**—Enable indirectly connected next hops.

no-indirect-next-hop—Explicitly disable indirect next hops.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Optimizing Route Reconvergence by Enabling Indirect Next Hops on the Packet Forwarding Engine*

Copyright © 2017, Juniper Networks, Inc.

install (Routing Options)

Syntax (install | no-install);

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit routing-options rib *routing-table-name* static (defaults | route)]
[edit routing-options static (defaults | route)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure whether Junos OS installs all static routes into the forwarding table. Even if you configure a route so it is not installed in the forwarding table, the route is still eligible to be exported from the routing table to other protocols.

Options
install—Explicitly install all static routes into the forwarding table. Include this statement when configuring an individual route in the **route** portion of the **static** statement to override a **no-install** option specified in the **defaults** portion of the statement.

no-install—Do not install the route into the forwarding table, even if it is the route with the lowest preference.

Default: **install**

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *static*

Copyright © 2017, Juniper Networks, Inc.

instance-export

Syntax instance-export [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Apply one or more policies to routes being exported from a routing instance.

Options ***policy-names***—Name of one or more export policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*

instance-import

Syntax instance-import [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Apply one or more policies to routes being imported into a routing instance.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Routing Policies, Firewall Filters, and Traffic Policers Feature Guide*

Copyright © 2017, Juniper Networks, Inc.

instance-type

Syntax instance-type virtual-router

Hierarchy Level [edit routing-instances]

Release Information Statement introduced in Junos OS Release 9.2 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Specify the type of routing instance.

Options **virtual-router**—Virtual router routing instance.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Using Virtual Routing Instances to Route Among VLANs on EX Series Switches*
- *Configuring Virtual Routing Instances (CLI Procedure)*
- Configuring Virtual Router Routing Instances on page 8

 Copyright © 2017, Juniper Networks, Inc.

interface (Multicast Static Routes)

Syntax

```
interface interface-names {
  disable;
  maximum-bandwidth bps;
  no-qos-adjust;
  reverse-oif-mapping {
    no-qos-adjust;
  }
  subscriber-leave-timer seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Enable multicast traffic on an interface.

By default, multicast packets are forwarded by enabling Protocol Independent Multicast (PIM) on an interface. PIM adds multicast routes into the routing table.

You can also configure multicast packets to be forwarded over a static route, such as a static route associated with an LSP next hop. Multicast packets are accepted on an interface and forwarded over a static route in the forwarding table. This is useful when you want to enable multicast traffic on a specific interface without configuring PIM on the interface.

You cannot enable multicast traffic on an interface and configure PIM on the same interface simultaneously.

Static routes must be configured before you can enable multicast on an interface. Configuring the **interface** statement alone does not install any routes into the routing table. This feature relies on the static route configuration.

Options

interface-names—Name of one or more interfaces on which to enable multicast traffic.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Defining Interface Bandwidth Maximums*
- *Example: Configuring Multicast with Subscriber VLANs*

Copyright © 2017, Juniper Networks, Inc.

interface (Routing Instances)

Syntax interface *interface-name*;

Hierarchy Level [edit routing-instances]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description For virtual router routing instances, configure an interface.

NOTE:

- You must configure only interfaces from the Node devices assigned to the network Node group. If you try to configure interfaces from the Node devices assigned to server Node groups, the configuration does not commit.
- You can configure an interface for one routing instance only.

Options ***interface-name***—Name of an interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Virtual Router Routing Instances on page 8

Copyright © 2017, Juniper Networks, Inc.

interface (Routing Options)

Syntax

```
interface interface-names {
  maximum-bandwidth bps;
  no-qos-adjust;
  reverse-oif-mapping {
    no-qos-adjust;
  }
  subscriber-leave-timer seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Enable multicast traffic on an interface.

TIP: You cannot enable multicast traffic on an interface by using the **routing-options multicast interface** statement and configure PIM on the interface.

Options

interface-name—Names of the physical or logical interface.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Defining Interface Bandwidth Maximums*
- *Example: Configuring Multicast with Subscriber VLANs*

Copyright © 2017, Juniper Networks, Inc.

interface-routes

Syntax

```
interface-routes {
  family (inet | inet6) {
    export {
      lan;
      point-to-point;
    }
  }
  rib-group group-name;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

Description

Associate a routing table group with the routing device's interfaces, and specify routing table groups into which interface routes are imported.

By default, IPv4 interface routes (also called direct routes) are imported into routing table **inet.0,** and IPv6 interface routes are imported into routing table **inet6.0**. If you are configuring alternate routing tables for use by some routing protocols, it might be necessary to import the interface routes into the alternate routing tables. To define the routing tables into which interface routes are imported, you create a routing table group and associate it with the routing device's interfaces.

To create the routing table groups, include the **passive** statement at the **[edit routing-options]** hierarchy level.

If you have configured a routing table, configure the OSPF primary instance at the **[edit protocols ospf]** hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group.

To export local routes, include the **export** statement.

To export LAN routes, include the **lan** option. To export point-to-point routes, include the **point-to-point** option.

Copyright © 2017, Juniper Networks, Inc.

Only local routes on point-to-point interfaces configured with a destination address are exportable.

Options

inet—Specify the IPv4 address family.

inet6—Specify the IPv6 address family.

lan—Export LAN routes.

point-to-point—Export point-to-point routes.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Importing Direct and Static Routes Into a Routing Instance*
- *Example: Configuring Multiple Routing Instances of OSPF*
- *passive*

local-address (Routing Options)

Syntax

local-address *address*;

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast backup-pe-group *group-name*],
[edit logical-systems *logical-system-name* routing-options multicast backup-pe-group *group-name*],
[edit routing-instances *routing-instance-name* routing-options multicast backup-pe-group *group-name*],
[edit routing-options multicast backup-pe-group *group-name*]

Release Information

Statement introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement added to the multicast hierarchy in Junos OS Release 13.2.

Description

Configure the address of the local PE for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.

Options

address—Address of local PEs in the backup group.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Ingress PE Redundancy*

Copyright © 2017, Juniper Networks, Inc.

martians

Syntax

```
martians {
  destination-prefix match-type <allow>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name*],
[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name*],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name*],
[edit routing-options],
[edit routing-options rib *routing-table-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure martian addresses.

Options

allow—(Optional) Explicitly allow a subset of a range of addresses that has been disallowed. The **allow** option is the only supported action.

destination-prefix—Destination route you are configuring:

- ***destination-prefix/prefix-length*—*destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.
- **default**—Default route to use when routing packets do not match a network or host in the routing table. This is equivalent to specifying the IP address **0.0.0.0/0**.

match-type—Criteria that the destination must match:

- **exact**—Exactly match the route's mask length.
- **longer**—The route's mask length is greater than the specified mask length.
- **orlonger**—The route's mask length is equal to or greater than the specified mask length.
- **through *destination-prefix***—The route matches the first prefix, the route matches the second prefix for the number of bits in the route, and the number of bits in the route is less than or equal to the number of bits in the second prefix.
- **upto *prefix-length***—The route's mask length falls between the two destination prefix lengths, inclusive.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring Class E Martian Addresses for Routing*

maximum-bandwidth

Syntax
maximum-bandwidth *bps*;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit routing-options multicast interface *interface-name*]

Release Information
Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description
Configure the multicast bandwidth for the interface.

Options
bps—Bandwidth rate, in bits per second, for the multicast interface.

Range: 0 through any amount of bandwidth

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Defining Interface Bandwidth Maximums*

Copyright © 2017, Juniper Networks, Inc.

maximum-paths

Syntax maximum-paths *path-limit* <log-interval *seconds*> <log-only | threshold *value*>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure a limit for the number of routes installed in a routing table based upon the route path.

NOTE: The maximum-paths statement is similar to the maximum-prefixes statement. The maximum-prefixes statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:

```
OSPF 10.10.10.0/24
ISIS 10.10.10.0/24
```

These are two routes, but only one destination (prefix). The maximum-paths limit applies the total number of routes (two). The maximum-prefixes limit applies to the total number of unique prefixes (one).

Options **log-interval *seconds***—(Optional) Minimum time interval (in seconds) between log messages.

Range: 5 through 86,400

log-only—(Optional) Sets the route limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.

path-limit—Maximum number of routes. If this limit is reached, a warning is triggered and additional routes are rejected.

Range: 1 through 4,294,967,295 ($2^{32} - 1$)

Default: No default

threshold *value*—(Optional) Percentage of the maximum number of routes that starts triggering a warning. You can configure a percentage of the ***path-limit*** value that starts triggering the warnings.

Range: 1 through 100

Copyright © 2017, Juniper Networks, Inc.

NOTE: When the number of routes reaches the **threshold** value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the ***path-limit*** value, then additional routes are rejected.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs*

Copyright © 2017, Juniper Networks, Inc.

maximum-prefixes

Syntax maximum-prefixes *prefix-limit* <log-interval *seconds*> <log-only | threshold *percentage*>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure a limit for the number of routes installed in a routing table based upon the route prefix.

Using a prefix limit, you can curtail the number of prefixes received from a CE router in a VPN. Prefix limits apply only to dynamic routing protocols and are not applicable to static or interface routes.

NOTE: The maximum-prefixes statement is similar to the maximum-paths statement. The maximum-prefixes statement limits the number of unique destinations in a routing instance. For example, suppose a routing instance has the following routes:

```
OSPF 10.10.10.0/24
ISIS 10.10.10.0/24
```

These are two routes, but only one destination (prefix). The maximum-paths limit applies the total number of routes (two). The maximum-prefixes limit applies to the total number of unique prefixes (one).

Options **log-interval *seconds***—(Optional) Minimum time interval (in seconds) between log messages.

Range: 5 through 86,400

log-only—(Optional) Sets the prefix limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.

prefix-limit—Maximum number of route prefixes. If this limit is reached, a warning is triggered and any additional routes are rejected.

Range: 1 through 4,294,967,295

Default: No default

Copyright © 2017, Juniper Networks, Inc.

threshold *value*—(Optional) Percentage of the maximum number of prefixes that starts triggering a warning. You can configure a percentage of the ***prefix-limit*** value that starts triggering the warnings.

Range: 1 through 100

NOTE: When the number of routes reaches the threshold value, routes are still installed into the routing table while warning messages are sent. When the number of routes reaches the *prefix-limit* value, then additional routes are rejected.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Limiting the Number of Paths and Prefixes Accepted from CE Routers in Layer 3 VPNs*

med-igp-update-interval

Syntax
med-igp-update-interval *minutes*;

Hierarchy Level
[edit routing-options]

Release Information
Statement introduced in Junos OS Release 9.0
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Configure a timer for how long to delay updates for the multiple exit discriminator (MED) path attribute for BGP groups and peers configured with the **metric-out igp *offset* delay-med-update** statement. The timer delays MED updates for the interval configured unless the MED is lower than the previously advertised attribute or another attribute associated with the route has changed or if the BGP peer is responding to a refresh route request.

Options
minutes—Interval to delay MED updates.

Range: 10 through 600

Default: 10 minutes

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Associating the MED Path Attribute with the IGP Metric and Delaying MED Updates*
- *metric-out*

Copyright © 2017, Juniper Networks, Inc.

metric (Aggregate, Generated, or Static Route)

Syntax (metric | metric2 | metric3 | metric4) *metric* <type *type*>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Specify the metric value for an aggregate, generated, or static route. You can specify up to four metric values, starting with **metric** (for the first metric value) and continuing with **metric2**, **metric3**, and **metric4**.

Options ***metric***—Metric value.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

type *type*—(Optional) Type of route.

When routes are exported to OSPF, type 1 routes are advertised in type 1 externals, and routes of any other type are advertised in type 2 externals. Note that if a qualified-next-hop metric value is configured, this value overrides the route metric.

Range: 1 through 16

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- aggregate on page 57
- generate on page 93
- *static*
- *Example: Summarizing Static Routes Through Route Aggregation*
- *Understanding Route Aggregation*

Copyright © 2017, Juniper Networks, Inc.

multicast (Routing Options)

Syntax

```
multicast {
  forwarding-cache {
    threshold suppress value <reuse value>;
  }
  interface interface-name {
    enable;
  }
  local-address address
  scope scope-name {
    interface [ interface-names ];
    prefix destination-prefix;
  }
  ssm-groups {
    address;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure generic multicast properties.

NOTE: You cannot apply a scoping policy to a specific routing instance. All scoping policies are applied to all routing instances. However, you can apply the scope statement to a specific routing instance.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Examples: Configuring Administrative Scoping*
- *Example: Configuring Source-Specific Multicast Groups with Any-Source Override*
- *Examples: Configuring the Multicast Forwarding Cache*
- *Multicast Protocols Feature Guide*
- (indirect-next-hop on page 99 | no-indirect-next-hop)

Copyright © 2017, Juniper Networks, Inc.

no-qos-adjust

Syntax no-qos-adjust;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name* reverse-oif-mapping],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name* reverse-oif-mapping],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name* reverse-oif-mapping],
[edit routing-options multicast interface *interface-name*],
[edit routing-options multicast interface *interface-name* reverse-oif-mapping]

Release Information
Statement introduced in Junos OS Release 9.5.
Statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement added to **[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*]**, **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*]**, and **[edit routing-options multicast interface *interface-name*]** hierarchy levels in Junos OS Release 9.6.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Disable hierarchical bandwidth adjustment for all subscriber interfaces that are identified by their MLD or IGMP request from a specific multicast interface.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring Multicast with Subscriber VLANs*

Copyright © 2017, Juniper Networks, Inc.

num-65-127-prefix

Syntax num-65-127-prefix *number*;

Hierarchy Level [edit chassis forwarding-options *profile-name*]

Release Information Statement introduced in Junos OS Release 13.2 for QFX Series switches.
Support for QFX5200 Series switches introduced in Junos OS Release 15.1x53-D30.

Description For the Unified Forwarding Table (UFT) feature, specify how much forwarding table memory to allocate for IPv6 entries with prefix lengths in the range of /65 through /127. The ability to allocate flexibly the memory for IPv6 entries with prefixes in this range extends the use of this memory space to accommodate the appropriate mix of longest-prefix match (LPM) entries that best suits your network. The LPM table stores IPv4 unicast prefixes, IPv6 prefixes with lengths equal to or less than 64, and IPv6 prefixes with lengths from 65 through 127. With this option, you can increase, decrease, or allocate no memory for IPv6 prefixes with lengths from 65 through 127, depending on which version of Junos OS you are using.

NOTE: This statement is supported only for the following forwarding table memory profiles: l2-profile-one, l2-profile-three, l2-profile-two, and l3-profile. Do not use this statement with the custom-profile or the lpm-profile statements.

NOTE: The values you can configure are different depending on the version of Junos OS you are using.

Options **number**—Specify a numerical value.

Range: (Junos OS Release 13.2x51-D10 only) 1 through 128. Each increment represents 16 IPv6 prefixes with lengths in the range of /65 through /127, for a total maximum of 2,058 prefixes (16 x 128 = 2,048).

Default: 1 (16 IPv6 prefixes with lengths in the range of /65 through /127).

Range: (Junos OS Release 13.2x51-D15 or later) 0 through 4. Each increment allocates memory for 1,000 IPv6 prefixes with lengths in the range of /65 through /127, for a maximum of 4,000 such IPv6 prefixes.

Default: 1 (1,000 IPv6 prefixes with lengths in the range of /65 through /127).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Unified Forwarding Table on page 39

Copyright © 2017, Juniper Networks, Inc.

options (Routing Options)

Syntax

```
options {
  syslog (level level | upto level level);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure the types of system logging messages sent about the routing protocols process to the system message logging file. These messages are also displayed on the system console. You can log messages at a particular level, or up to and including a particular level.

Options

level *level*—Severity of the message. It can be one or more of the following levels, in order of decreasing urgency:

- **alert**—Conditions that should be corrected immediately, such as a corrupted system database.
- **critical**—Critical conditions, such as hard drive errors.
- **debug**—Software debugging messages.
- **emergency**—Panic or other conditions that cause the system to become unusable.
- **error**—Standard error conditions.
- **info**—Informational messages.
- **notice**—Conditions that are not error conditions, but might warrant special handling.
- **warning**—System warning messages.

upto level *level*—Log all messages up to a particular level.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **syslog** in the *Junos OS Administration Library*

Copyright © 2017, Juniper Networks, Inc.

pim-to-igmp-proxy

Syntax

```
pim-to-igmp-proxy {
  upstream-interface [ interface-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description

Configure the rendezvous point (RP) routing device that resides between a customer edge-facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Internet Group Management Protocol (IGMP) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the **pim-to-igmp-proxy** statement enables you to use IGMP to forward IPv4 multicast traffic across the PIM sparse mode domains.

The **pim-to-igmp-proxy** statement is not supported for routing instances.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring PIM-to-IGMP Message Translation*

Copyright © 2017, Juniper Networks, Inc.

pim-to-mld-proxy

Syntax

```
pim-to-mld-proxy {
  upstream-interface [ interface-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description

Configure the rendezvous point (RP) routing device that resides between a customer edge–facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Multicast Listener Discovery (MLD) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the **pim-to-mld-proxy** statement enables you to use MLD to forward IPv6 multicast traffic across the PIM sparse mode domains.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring PIM-to-MLD Message Translation*

Copyright © 2017, Juniper Networks, Inc.

policy (Aggregate and Generated Routes)

Syntax policy *policy-name*;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (aggregate | generate) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)],
[edit routing-options (aggregate | generate) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate) (defaults | route)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Associate a routing policy when configuring an aggregate or generated route's destination prefix in the **routes** part of the **aggregate** or **generate** statement. This provides the equivalent of an import routing policy filter for the destination prefix. That is, each potential contributor to an aggregate route, along with any aggregate options, is passed through the policy filter. The policy then can accept or reject the route as a contributor to the aggregate route.

If the contributor is accepted, the policy can modify the default preferences. The contributor with the numerically smallest prefix becomes the most preferred, or *primary*, contributor. A rejected contributor still can contribute to a less specific aggregate route. If you do not specify a policy filter, all candidate routes contribute to an aggregate route.

The following algorithm is used to compare two generated contributing routes in order to determine which one is the primary or preferred contributor:

1. Compare the protocol's **preference** of the contributing routes. The lower the preference, the better the route. This is similar to the comparison that is done while determining the best route for the routing table.
2. Compare the protocol's **preference2** of the contributing routes. The lower **preference2** value is better. If only one route has **preference2**, then this route is preferred.
3. The preference values are the same. Proceed with a numerical comparison of the prefixes' values.
   a. The primary contributor is the numerically smallest prefix value.

Copyright © 2017, Juniper Networks, Inc.

b. If the two prefixes are numerically equal, the primary contributor is the route that has the smallest prefix length value.

At this point, the two routes are the same. The primary contributor does not change. An additional next hop is available for the existing primary contributor.

A rejected contributor still can contribute to less specific generated route. If you do not specify a policy filter, all candidate routes contribute to a generated route.

Options ***policy-name***—Name of a routing policy.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Summarizing Static Routes Through Route Aggregation*
- aggregate on page 57
- generate on page 93

policy (Flow Maps)

Syntax policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-options multicast flow-map *flow-map-name*]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure a flow map policy.

Options ***policy-names***—Name of one or more policies for flow mapping.

Required Privilege Level routing—To view this statement in the configuration.

Copyright © 2017, Juniper Networks, Inc.

policy-options

Syntax

```
policy-options
  application-maps application-map-name {
    application application-name {
      code-points [ aliases ] [ bit-patterns ];
    }
  }
  policy-statement policy-name {
    term term-name {
      from {
        family family-name;
        match-conditions;
        policy subroutine-policy-name;
        prefix-list prefix-list-name;
        prefix-list-filter prefix-list-name match-type <actions>;
        route-filter destination-prefix match-type <actions>;
        source-address-filter source-prefix match-type <actions>;
      }
      to {
        match-conditions;
        policy subroutine-policy-name;
      }
      then actions;
    }
  }
```

Hierarchy Level [edit]

Release Information
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.1 for the EX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure options such as application maps for DCBX application protocol exchange and policy statements.

Required Privilege Level
storage—To view this statement in the configuration.
storage-control—To add this statement to the configuration.

Related Documentation

- *Defining an Application for DCBX Application Protocol TLV Exchange*
- *Example: Configuring DCBX Application Protocol TLV Exchange*
- *Example: Configuring DCBX to Support an iSCSI Application*
- *Understanding DCBX Application Protocol TLV Exchange*
- *Understanding DCBX Application Protocol TLV Exchange on EX Series Switches*

Copyright © 2017, Juniper Networks, Inc.

policy-statement

Syntax

```
policy-statement policy-name {
  term term-name {
    from {
      family family-name;
      match-conditions;
      policy subroutine-policy-name;
      prefix-list prefix-list-name;
      prefix-list-filter prefix-list-name match-type <actions>;
      protocol protocol-name;
      route-filter destination-prefix match-type <actions>;
      source-address-filter source-prefix match-type <actions>;
      traffic-engineering;
    }
    to {
      match-conditions;
      policy subroutine-policy-name;
    }
    then actions;
  }
  then {
    no-entropy-label-capability;
  }
}
```

Hierarchy Level

[edit dynamic policy-options],
[edit logical-systems *logical-system-name* policy-options],
[edit policy-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for configuration in the dynamic database introduced in Junos OS Release 9.5.
Support for configuration in the dynamic database introduced in Junos OS Release 9.5 for EX Series switches.
inet-mdt option introduced in Junos OS Release 10.0R2.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
route-target option introduced in Junos OS Release 12.2.
Statement introduced in Junos OS 14.1X53-D20 for the OCX Series.
protocol and **traffic-engineering** options introduced in Junos OS Release 14.2.
no-entropy-label-capability option introduced in Junos OS Release 15.1.

Description

Define a routing policy, including subroutine policies.

A *term* is a named structure in which match conditions and actions are defined. Routing policies are made up of one or more terms. Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each term contains a set of match conditions and a set of actions:

Copyright © 2017, Juniper Networks, Inc.

- Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route.
- Actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

Generally, a router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of **accept** or **reject** is taken. If none of the terms in the policy match the route, the router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

If none of the match conditions of each term evaluates to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either **accept** or **reject**) that overrides any action intrinsic to the protocol.

The order of match conditions in a term is not relevant, because a route must match all match conditions in a term for an action to be taken.

To list the routing policies under the **[edit policy-options]** hierarchy level by **policy-statement** ***policy-name*** in alphabetical order, enter the **show policy-options** configuration command.

Copyright © 2017, Juniper Networks, Inc.

Options

actions—(Optional) One or more actions to take if the conditions match. The actions are described in *Configuring Flow Control Actions*.

family *family-name*—(Optional) Specify an address family protocol. Specify **inet** for IPv4. Specify **inet6** for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify **iso**. For IPv4 multicast VPN traffic, specify **inet-mvpn**. For IPv6 multicast VPN traffic, specify **inet6-mvpn**. For multicast-distribution-tree (MDT) IPv4 traffic, specify **inet-mdt**. For BGP route target VPN traffic, specify **route-target**. For traffic engineering, specify **traffic-engineering**.

NOTE: When **family** is not specified, the routing device or routing instance uses the address family or families carried by BGP. If multiprotocol BGP (MP-BGP) is enabled, the policy defaults to the protocol family or families carried in the network layer reachability information (NLRI) as configured in the *family* statement for BGP. If MP-BGP is not enabled, the policy uses the default BGP address family unicast IPv4.

from—(Optional) Match a route based on its source address.

match-conditions—(Optional in **from** statement; required in **to** statement) One or more conditions to use to make a match. The qualifiers are described in *Routing Policy Match Conditions*.

no-entropy-label-capability—(Optional) Disable the entropy label capability advertisement at egress or transit routes specified in the policy.

policy *subroutine-policy-name*—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”). Policy names cannot take the form **__.*-internal__**, as this form is reserved. For information about how to configure subroutines, see *Understanding Policy Subroutines in Routing Policy Match Conditions*.

policy-name—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

prefix-list *prefix-list-name*—Name of a list of IPv4 or IPv6 prefixes.

prefix-list-filter *prefix-list-name*—Name of a prefix list to evaluate using qualifiers; ***match-type*** is the type of match (see *Configuring Prefix List Filters*), and ***actions*** is the action to take if the prefixes match.

protocol *protocol-name*—Name of the protocol used to control traffic engineering database import at the originating point.

route-filter *destination-prefix match-type <actions>*—(Optional) List of routes on which to perform an immediate match; ***destination-prefix*** is the IPv4 or IPv6 route prefix to

Copyright © 2017, Juniper Networks, Inc.

match, ***match-type*** is the type of match (see *Configuring Route Lists*), and ***actions*** is the action to take if the ***destination-prefix*** matches.

source-address-filter ***source-prefix match-type*** **<*actions*>**—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. ***source-prefix*** is the IPv4 or IPv6 route prefix to match, ***match-type*** is the type of match (see *Configuring Route Lists*), and ***actions*** is the action to take if the ***source-prefix*** matches.

term ***term-name***—Name that identifies the term. The term name must be unique in the policy. It can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" "). A policy statement can include multiple terms. We recommend that you name all terms. However, you do have the option to include an unnamed term which must be the final term in the policy. To configure an unnamed term, omit the **term** statement when defining match conditions and actions.

to—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

then—(Optional) Actions to take on matching routes. The actions are described in *Configuring Flow Control Actions* and *Configuring Actions That Manipulate Route Characteristics*.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *dynamic-db*

Copyright © 2017, Juniper Networks, Inc.

ppm

Syntax

```
ppm {
  no-delegate-processing;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information Statement introduced in Junos OS Release 9.4.
Statement introduced in Junos OS Release 10.2 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description (M120, M320, MX Series, T Series, TX Matrix routers, M7i and M10i routers with Enhanced CFEB [CFEB-E], EX Series switches, and QFX Series only) Disable distributed periodic packet management (PPM) to the Packet Forwarding Engine (on routers), to access ports (on EX3200 and EX4200 switches, and QFX Series), or to line cards (on EX6200 and EX8200 switches).

After you disable PPM, PPM processing continues to run on the Routing Engine.

In Junos OS Release 8.2, PPM was moved from the Routing Engine to the Packet Forwarding Engine, access ports, or line cards. The **no-delegate-processing** statement disables the default behavior and restores the legacy behavior.

Default Distributed PPM processing is enabled for all protocols that use PPM.

Options **no-delegate-processing**—Disable PPM to the Packet Forwarding Engine, access ports, or line cards. Distributed PPM is enabled by default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Ensuring That Distributed ppm Is Not Disabled*

Copyright © 2017, Juniper Networks, Inc.

ppm (Ethernet Switching)

Syntax

```
ppm {
  centralized;
}
```

Hierarchy Level [edit protocols lacp]

Release Information

Statement introduced in Junos OS Release 9.4 for MX Series routers.
Statement introduced in Junos OS Release 10.2 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.1 for T Series devices.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure PPM processing options for Link Aggregation Control Protocol (LACP) packets.

This command configures the PPM processing options for LACP packets only. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the **no-delegate-processing** configuration statement in the **[edit routing-options** ppm**]** statement hierarchy.

Default Distributed PPM processing is enabled for all packets that use PPM.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Configuring Distributed Periodic Packet Management on an EX Series Switch (CLI Procedure)*
- Configuring Distributed Periodic Packet Management on page 9

Copyright © 2017, Juniper Networks, Inc.

preference (Routing Options)

Syntax

```
preference {
  metric-value;
  <type metric_type>
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Preference value for a static, aggregate, or generated route. You also can specify a secondary preference value, as well as color values, which are even finer-grained preference values.

You can specify a primary route preference (by including the **preference** statement in the configuration), and a secondary preference that is used as a tiebreaker (by including the **preference2** statement). You can also mark route preferences with additional route tiebreaker information by specifying a color and a tiebreaker color (by including the **color** and **color2** statements in the configuration).

If the Junos OS routing table contains a dynamic route to a destination that has a better (lower) preference value than the static, aggregate, or generated route, the dynamic route is chosen as the active route and is installed in the forwarding table.

Options

metric_value—The metric value for an aggregate, a generated, or a static route to determine the best route among multiple routes to a destination

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: 5 (for static routes), 130 (for aggregate and generated routes)

type *metric_type*—(Optional) External metric type for routes exported by OSPF. When routes are exported to OSPF, type 1 routes are advertised in type 1 externals, and routes of any other type are advertised in type 2 externals. Note that if a qualified-next-hop metric value is configured, this value overrides the route metric.

Copyright © 2017, Juniper Networks, Inc.

Range: 1 through 16

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Summarizing Static Routes Through Route Aggregation*
- aggregate on page 57
- generate on page 93
- *static*
- *color*

prefix

Syntax
prefix *destination-prefix*;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast scope *scope-name*],
[edit logical-systems *logical-system-name* routing-options multicast scope *scope-name*],
[edit routing-instances *routing-instance-name* routing-options multicast scope *scope-name*],
[edit routing-options multicast scope *scope-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Configure the prefix for multicast scopes.

Options
destination-prefix—Address range for the multicast scope.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Examples: Configuring Administrative Scoping*
- *Example: Creating a Named Scope for Multicast Scoping*
- *multicast*

Copyright © 2017, Juniper Networks, Inc.

prefix-65-127-disable

Syntax prefix-65-127-disable;

Hierarchy Level [edit chassis forwarding-options lpm-profile]

Release Information Statement introduced in Junos OS Release 13.2X51-D15 for QFX Series switches.
Support introduced in Junos OS Release 15.1x53-D30 for QFX5200 Series switches.

Description For the Unified Forwarding Table (UFT) feature, specify not to allocate any memory for IPv6 prefixes with lengths in the range /65 through /127 for longest-prefix-match (LPM) entries. Doing so increases the memory available for LPM entries for IPv4 unicast prefixes and IPv6 prefixes with lengths equal to or less than 64. The maximum default value for LPM entries is 16,000 IPv6 prefixes of all lengths.

In an environment where the switch is being used in the core of the network, for example, it might not need to store IPv6 prefixes with lengths in the range /65 through /127. IPv6 prefixes of this type are not typically used in the core.

NOTE: On the QFX3500, QFX3600, and QFX5100 switches, when you configure this statement, the maximum number of LPM IPv6 entries with prefix lengths equal to or less than 64 increases to 128,000. On the QFX5200 switch, when you configure this statement, the maximum number of IPv6 entries with prefix lengths equal to or less 64 that are allocated in the LPM table increases to 98,0000.

NOTE: This statement is supported only with the lpm-profile. No other profile is supported.

Options **None**—This statement has no options.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Unified Forwarding Table on page 39
- *Understanding the Unified Forwarding Table*

Copyright © 2017, Juniper Networks, Inc.

protocols

Syntax

```
protocols {
  bgp {
    ... bgp-configuration ...
  }
  isis {
    ... isis-configuration ...
  }
  ldp {
    ... ldp-configuration ...
  }
  mpls {
    ... mpls -configuration ...
  }
  msdp {
    ... msdp-configuration ...
  }
  mstp {
    ... mstp-configuration ...
  }
  ospf {
    domain-id domain-id;
    domain-vpn-tag number;
    route-type-community (iana | vendor);
    traffic-engineering {
      <advertise-unnumbered-interfaces>;
      <credibility-protocol-preference>;
      ignore-lsp-metrics;
      multicast-rpf-routes;
      no-topology;
      shortcuts {
        lsp-metric-into-summary;
      }
    }
    ... ospf-configuration ...
  }
  ospf3 {
    domain-id domain-id;
    domain-vpn-tag number;
    route-type-community (iana | vendor);
    traffic-engineering {
      <advertise-unnumbered-interfaces>;
      <credibility-protocol-preference>;
      ignore-lsp-metrics;
      multicast-rpf-routes;
      no-topology;
      shortcuts {
        lsp-metric-into-summary;
      }
    }
    ... ospf3-configuration ...
  }
  pim {
```

Copyright © 2017, Juniper Networks, Inc.

```
    ... pim-configuration ...
  }
  rip {
    ... rip-configuration ...
  }
  ripng {
    ... ripng-configuration ...
  }
  rstp {
    rstp-configuration;
  }
  rsvp{
    ... rsvp-configuration ...
  }
  vstp {
    vstp configuration;
  }
  vpls {
    vpls configuration;
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*], [edit routing-instances *routing-instance-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Support for RIPng introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 11.1 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
mpls and **rsvp** options added in Junos OS Release 15.1.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Specify the protocol for a routing instance. You can configure multiple instances of many protocol types. Not all protocols are supported on the switches. See the switch CLI.

Copyright © 2017, Juniper Networks, Inc.

Options

bgp—Specify BGP as the protocol for a routing instance.

isis—Specify IS-IS as the protocol for a routing instance.

ldp—Specify LDP as the protocol for a routing instance or for a virtual router instance.

l2vpn—Specify Layer 2 VPN as the protocol for a routing instance.

mpls—Specify MPLS as the protocol for a routing instance.

msdp—Specify the Multicast Source Discovery Protocol (MSDP) for a routing instance.

mstp—Specify the Multiple Spanning Tree Protocol (MSTP) for a virtual switch routing instance.

ospf—Specify OSPF as the protocol for a routing instance.

ospf3—Specify OSPF version 3 (OSPFv3) as the protocol for a routing instance.

NOTE: **OSPFv3 supports the no-forwarding, virtual-router, and vrf routing instance types only.**

pim—Specify the Protocol Independent Multicast (PIM) protocol for a routing instance.

rip—Specify RIP as the protocol for a routing instance.

ripng—Specify RIP next generation (RIPng) as the protocol for a routing instance.

rstp—Specify the Rapid Spanning Tree Protocol (RSTP) for a virtual switch routing instance.

rsvp—Specify the RSVP for a routing instance.

vstp—Specify the VLAN Spanning Tree Protocol (VSTP) for a virtual switch routing instance.

vpls—Specify VPLS as the protocol for a routing instance.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Multiple Routing Instances of OSPF*

Copyright © 2017, Juniper Networks, Inc.

qualified-next-hop (Static Routes)

Syntax

```
qualified-next-hop (address | interface-name) {
  bfd-liveness-detection {
    authentication {
      algorithm (keyed-md5 | keyed-sha-1 | meticulous-keyed-md5 | meticulous-keyed-sha-1 |
        simple-password);
      key-chain key-chain-name;
      loose-check;
    }
    detection-time {
      threshold milliseconds;
    }
    holddown-interval milliseconds;
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
      minimum-interval milliseconds;
      threshold milliseconds;
    }
    version (1 | automatic);
  }
  interface interface-name;
  metric metric;
  preference preference;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-options rib inet6.0 static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-options static route *destination-prefix*],
[edit routing-instances *routing-instance-name* routing-options static route *destination-prefix*],
[edit routing-options rib inet6.0 static route *destination-prefix*],
[edit routing-options static route *destination-prefix*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure a static route with multiple possible next hops, each of which can have its own preference value, IGP metric that is used when the route is exported into an IGP, and Bidirectional Forwarding Detection (BFD) settings. If multiple links are operational, the one with the most preferred next hop is used. The most preferred next hop is the one with the lowest preference value.

Options

address—IPv4, IPv6, or ISO network address of the next hop.

Copyright © 2017, Juniper Networks, Inc.

interface-name—Name of the interface on which to configure an independent metric or preference for a static route. To configure an unnumbered interface as the next-hop interface for a static route, specify **qualified-next-hop *interface-name***, where ***interface-name*** is the name of the IPv4 or IPv6 unnumbered interface.

NOTE: For an Ethernet interface to be configured as the qualified next hop for a static route, it must be an unnumbered interface.

To configure an Ethernet interface as an unnumbered interface, configure the **unnumbered-address** *<interface-name>* statement at the **[edit interfaces** *<interface-name>* **unit** *<logical-unit-number>* **family** *<family-name>***]** hierarchy level as described in *Configuring an Unnumbered Interface*.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Enabling BFD on Qualified Next Hops in Static Routes for Route Selection on page 22

Copyright © 2017, Juniper Networks, Inc.

readvertise

Syntax (readvertise | no-readvertise);

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit routing-options rib *routing-table-name* static (defaults | route)],
[edit routing-options static (defaults | route)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure whether static routes are eligible to be readvertised by routing protocols:

Default Static routes are eligible to be readvertised (that is, exported from the routing table into dynamic routing protocols) if a policy to do so is configured. To mark an IPv4 static route as being ineligible for readvertisement, include the **no-readvertise** statement.

Options
readvertise—Readvertise static routes. Include the **readvertise** statement when configuring an individual route in the **route** portion of the **static** statement to override a **no-readvertise** option specified in the **defaults** portion of the statement.

no-readvertise—Mark a static route as being ineligible for readvertisement. Include the **no-readvertise** option when configuring the route.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Preventing a Static Route from Being Readvertised*
- *Understanding Static Route Control in Routing and Forwarding Tables*
- *static*

Copyright © 2017, Juniper Networks, Inc.

redundant-sources

Syntax redundant-sources [*addresses*];

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-options multicast flow-map *flow-map-name*]

Release Information
Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Configure a list of redundant sources for multicast flows defined by a flow map.

Options ***addresses***—List of IPv4 or IPv6 addresses for use as redundant (backup) sources for multicast flows defined by a flow map.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring a Multicast Flow Map*

Copyright © 2017, Juniper Networks, Inc.

resolution

Syntax

```
resolution {
  rib routing-table-name {
    import [ policy-names ];
    resolution-ribs [ routing-table-names ];
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure the router to perform custom route resolution on protocol next hops of routes in a certain routing table. The protocol next hop is used to determine the forwarding next-hop.

For example, you might want to direct **inet.2** route resolution to use topology routing tables **:red.inet.0** and **:blue.inet.0** for protocol next-hop IP address lookups. Or you might want to direct bgp.l3vpn.0 to use the information in inet.0 to resolve routes, thus overriding the default behavior, which is to use inet.3.

You can specify up to two routing tables in the resolution-ribs statement. The route resolution scheme first checks the first-listed routing table for the protocol next-hop address. If the address is found, it uses this entry. If it is not found, the resolution scheme checks second-listed routing table. Hence, only one routing table is used for each protocol nexthop address. For example, if you configure **resolution rib bgp.l3vpn.0 resolution-ribs [inet.0 inet.3]**, inet.0 is checked first and then inet.3 is checked.

NOTE: Customizing route resolution might cause the routing protocol process (rpd) to consume more memory resources than it ordinarily would. When you customize route resolution, we recommend that you check the memory resources by running the **show system processes** and the **show task memory** commands. For more information, see *Routing Protocol Process Memory FAQs*.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- *Example: Configuring Route Resolution on PE Routers*
- *Example: Configuring Route Resolution on Route Reflectors*
- *Understanding Multitopology Routing in Conjunction with PIM*
- *Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths*

resolution-ribs

Syntax resolution-ribs [*routing-table-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options resolution rib],
[edit logical-systems *logical-system-name* routing-options resolution rib],
[edit routing-instances *routing-instance-name* routing-options resolution rib],
[edit routing-options resolution rib]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Specify one or more routing tables to use for route resolution.

This statement enables you to override the default routing tables that Junos OS uses for route resolution. For example, suppose that the resolution routing table is **inet.3**, but you want to allow fallback resolution through **inet.0**. One example use case is overriding the **bgp.rtarget.0** (**family route-target**) routing table resolution from using only **inet.3** to using both **inet.3** and **inet.0**.

Options ***routing-table-names***—Name of one or more routing tables.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Route Resolution on PE Routers*
- *Example: Configuring Multitopology Routing to Provide Redundancy for Multicast Traffic over Separate Network Paths*
- *Understanding Multitopology Routing in Conjunction with PIM*

Copyright © 2017, Juniper Networks, Inc.

resolve

Syntax resolve;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit routing-options rib *routing-table-name* static (defaults | route)],
[edit routing-options static (defaults | route)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Statically configure routes to be resolved to a next hop that is not directly connected. The route is resolved through the **inet.0** and **inet.3** routing tables.

NOTE: You cannot configure both **resolve** and **retain** options for a statically configured route because resolved next hops cannot be retained.

Default Static routes can point only to a directly connected next hop.

TIP: We recommend configuring the **no-resolve** option for individual routes to override default configuration.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *static*

Copyright © 2017, Juniper Networks, Inc.

restart-duration (Routing Options)

Syntax restart-duration *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options graceful-restart],
[edit logical-systems *logical-system-name* routing-options graceful-restart],
[edit routing-instances *routing-instance-name* routing-options graceful-restart],
[edit routing-options graceful-restart]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Configure the restart timer for graceful restart.

Options ***seconds***—Configure the time period for the restart to last.

Range: 120 through 900 seconds

Default: 300 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Junos OS High Availability Library for Routing Devices*

Copyright © 2017, Juniper Networks, Inc.

retain

Syntax (no-retain | retain);

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit logical-systems *logical-system-name* routing-options static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options static (defaults | route)],
[edit routing-options rib *routing-table-name* static (defaults | route)],
[edit routing-options static (defaults | route)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure statically configured routes to be deleted from or retained in the forwarding table when the routing protocol process shuts down normally:

NOTE: You cannot configure both **retain** and **resolve** options for a statically configured route because resolved next hops cannot be retained.

Default Statically configured routes are deleted from the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.

TIP: We recommend configuring the **no-retain** option for individual routes to override default configuration.

Options

no-retain—Delete statically configured routes from the forwarding table when the routing protocol process shuts down normally. To explicitly specify that routes be deleted from the forwarding table, include the **no-retain** statement. Include this statement when configuring an individual route in the **route** portion of the **static** statement to override a **retain** option specified in the **defaults** portion of the statement.

retain—Have a static route remain in the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *static*

reverse-oif-mapping

Syntax

```
reverse-oif-mapping {
  no-qos-adjust;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit routing-options multicast interface *interface-name*]

Release Information
Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 9.2 for EX Series switches.
The **no-qos-adjust** statement added in Junos OS Release 9.5.
The **no-qos-adjust** statement introduced in Junos OS Release 9.5 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description
Enable the routing device to identify a subscriber VLAN or interface based on an IGMP or MLD request it receives over the multicast VLAN.

The remaining statement is explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring Multicast with Subscriber VLANs*

Copyright © 2017, Juniper Networks, Inc.

rpf-check-policy (Routing Options RPF)

Syntax rpf-check-policy [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information
Statement introduced in Junos OS Release 8.1.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Apply policies for disabling RPF checks on arriving multicast packets. The policies must be correctly configured.

Options ***policy-names***—Name of one or more multicast RPF check policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring RPF Policies*

Copyright © 2017, Juniper Networks, Inc.

rib (General)

Syntax

```
rib routing-table-name {
  aggregate {
    defaults {
      ... aggregate-options ...
    }
    route destination-prefix {
      policy policy-name;
      ... aggregate-options ...
    }
    generate {
      defaults {
        generate-options;
      }
      route destination-prefix {
        policy policy-name;
        generate-options;
      }
    }
    martians {
      destination-prefix match-type <allow>;
    }
  }
  static {
    defaults {
      static-options;
    }
    rib-group group-name;
    route destination-prefix {
      next-hop;
      static-options;
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Create a routing table.

Explicitly creating a routing table with ***routing-table-name*** is optional if you are not adding any static, martian, aggregate, or generated routes to the routing table and if you also are creating a routing table group.

Copyright © 2017, Juniper Networks, Inc.

NOTE: The IPv4 multicast routing table (**inet.1**) and the IPv6 multicast routing table (**inet6.1**) are not supported for this statement.

Default If you do not specify a routing table name with the ***routing-table-name*** option, the software uses the default routing tables, which are **inet.0** for unicast routes and **inet.1** for the multicast cache.

Options ***routing-table-name***—Name of the routing table, in the following format: ***protocol*** [***.identifier***].

In a routing instance, the routing table name must include the routing instance name. For example, if the routing instance name is **link0**, the routing table name might be **link0.inet6.0**.

- ***protocol*** is the protocol family. It can be **inet6** for the IPv6 family, **inet** for the IPv4 family, **iso** for the ISO protocol family, or ***instance-name*.iso.0** for an ISO routing instance.
- ***identifier*** is a positive integer that specifies the instance of the routing table.

Default: **inet.0**

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Creating Routing Tables*
- *passive*

Copyright © 2017, Juniper Networks, Inc.

rib (Route Resolution)

Syntax

```
rib routing-table-name {
  import [ policy-names ];
  resolution-ribs [ routing-table-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options resolution],
[edit logical-systems *logical-system-name* routing-options resolution],
[edit routing-instances *routing-instance-name* routing-options resolution],
[edit routing-options resolution]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Specify a routing table name for route resolution.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Route Resolution on PE Routers*

Copyright © 2017, Juniper Networks, Inc.

rib-group (Routing Options)

Syntax rib-group *group-name*;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options interface-routes],
[edit logical-systems *logical-system-name* routing-options interface-routes],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static],
[edit logical-systems *logical-system-name* routing-options static],
[edit routing-instances *routing-instance-name* routing-options interface-routes],
[edit routing-options interface-routes],
[edit routing-options rib *routing-table-name* static],
[edit routing-options static]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure which routing table groups interface routes are imported into.

Options
group-name—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. It generally does not make sense to specify more than a single routing table group.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Importing Direct and Static Routes Into a Routing Instance*
- *Example: Exporting Specific Routes from One Routing Table Into Another Routing Table*
- interface-routes on page 106
- rib-groups on page 151

Copyright © 2017, Juniper Networks, Inc.

rib-groups

Syntax

```
rib-groups {
  group-name {
    export-rib group-name;
    import-policy [ policy-names ];
    import-rib [ group-names ];
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Group one or more routing tables to form a routing table group. A routing protocol can import routes into all the routing tables in the group and can export routes from a single routing table.

Each routing table group must contain one or more routing tables that Junos OS uses when importing routes (specified in the **import-rib** statement) and optionally can contain one routing table group that Junos OS uses when exporting routes to the routing protocols (specified in the **export-rib** statement).

The first routing table you specify is the *primary routing table*, and any additional routing tables are the *secondary routing tables*.

The primary routing table determines the address family of the routing table group. To configure an IP version 4 (IPv4) routing table group, specify **inet.0** as the primary routing table. To configure an IP version 6 (IPv6) routing table group, specify **inet6.0** as the primary routing table. If you configure an IPv6 routing table group, the primary and all secondary routing tables must be IPv6 routing tables (**inet6.*x***).

In Junos OS Release 9.5 and later, you can include both IPv4 and IPv6 routing tables in an IPv4 import routing table group using the **import-rib** statement. In releases prior to Junos OS Release 9.5, you can only include either IPv4 or IPv6 routing tables in the same **import-rib** statement. The ability to configure an import routing table group with both IPv4 and IPv6 routing tables enables you, for example, to populate the **inet6.3** routing table with IPv6 addresses that are compatible with IPv4. Specify **inet.0** as the primary routing table, and specify **inet6.3** as a secondary routing table.

NOTE: On EX Series switches, only dynamically learned routes can be imported from one routing table group to another.

Copyright © 2017, Juniper Networks, Inc.

NOTE: If you configure an import routing table group that includes both IPv4 and IPv6 routing tables, any corresponding export routing table group must include only IPv4 routing tables.

If you have configured a routing table, configure the OSPF primary instance at the **[edit protocols ospf]** hierarchy level with the statements needed for your network so that routes are installed in **inet.0** and in the forwarding table. Make sure to include the routing table group. For more information, see *Example: Configuring Multiple Routing Instances of OSPF*.

After specifying the routing table from which to import routes, you can apply one or more policies to control which routes are installed in the routing table group. To apply a policy to routes being imported into the routing table group, include the **import-policy** statement.

Options

group-name—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Exporting Specific Routes from One Routing Table Into Another Routing Table*
- rib-group on page 150

Copyright © 2017, Juniper Networks, Inc.

route-distinguisher-id

Syntax route-distinguisher-id *ip-address*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Automatically assign a route distinguisher to the routing instance.

If you configure the **route-distinguisher** statement in addition to the **route-distinguisher-id** statement, the value configured for **route-distinguisher** supersedes the value generated from **route-distinguisher-id**.

NOTE: To avoid a conflict in the two route distinguisher values, it is recommended to ensure that the first half of the route distinguisher obtained by configuring the **route-distinguisher** statement is different from the first half of the route distinguisher obtained by configuring the **route-distinguisher-id** statement.

Options ***ip-address***—Address for routing instance.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring BGP Route Target Filtering for VPNs*
- *Configuring Routing Instances on PE Routers in VPNs*

Copyright © 2017, Juniper Networks, Inc.

route-record

Syntax route-record;

Hierarchy Level [edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Export the AS path and routing information to the traffic sampling process.

Before you can perform flow aggregation, the routing protocol process must export the AS path and routing information to the sampling process.

NOTE: Starting with Junos OS Release 15.1, when you commit a minor configuration change, the routing protocol process sends only AS paths that are active routes to the FPCs. Not all known AS paths are sent to the FPC, thereby considerably reducing the memory and CPU usage, resulting in a faster route record database update.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Enabling Flow Aggregation*
- *Junos OS Services Interfaces Library for Routing Devices*

Copyright © 2017, Juniper Networks, Inc.

router-id

Syntax router-id *address*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Specify the routing device's IP address.

The router identifier is used by BGP and OSPF to identify the routing device from which a packet originated. The router identifier usually is the IP address of the local routing device. If you do not configure a router identifier, the IP address of the first interface to come online is used. This is usually the loopback interface. Otherwise, the first hardware interface with an IP address is used.

NOTE: We strongly recommend that you configure the router identifier under the **[edit routing-options]** hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.

You must configure a router-id in order for BGP and OSPF to function in a routing instance. Use the **show route instance detail** command to display the router-id value for a routing instance. If the router-id is **0.0.0.0**, then the routing instance has no router-id.

For more information about the router identifier in OSPF, see *Example: Configuring an OSPF Router Identifier*.

NOTE: If you run OSPF for IPv6 or BGP for IPv6 in a routing instance, you must configure an IPv4 router identifier (**router-id**) in the routing instance itself. In other words, the IPv4 **router-id** in the main routing instance is not inherited by other routing instances. Even if you run *only* IPv6 OSPF or BGP in a routing instance, the IPv4 **router-id** must be configured because OSPF and BGP, even when used exclusively with IPv6, use the IPv4 **router-id** for handshaking. If you do not configure the IPv4 **router-id** in the IPv6 OSPF or BGP routing instance, then the IPv6 protocols will use invalid IPv4 address **0.0.0.0** and the adjacencies and connections will fail.

Copyright © 2017, Juniper Networks, Inc.

Options ***address***—IP address of the routing device.

Default: Address of the first interface encountered by Junos OS

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Examples: Configuring External BGP Peering*
- *Examples: Configuring Internal BGP Peering*

routing-instances

Syntax

```
routing-instances routing-instance-name {
  description;
  instance-type virtual-router;
  interface interface-name;
  protocols;
  routing-options
}
```

Hierarchy Level [edit]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description (QFabric switches only) Configure a virtual router routing instance.

Options ***routing-instance-name***—Name of this routing instance.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Virtual Router Routing Instances on page 8

Copyright © 2017, Juniper Networks, Inc.

routing-options

Syntax routing-options { ... }

Hierarchy Level [edit],
[edit routing-instances *routing-instance-name*]

Release Information Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure protocol-independent routing properties.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Overview of Routing Options on page 3
- Understanding Distributed Periodic Packet Management on page 9
- *Example: Configuring SSM Maps for Different Groups to Different Sources*

scope

Syntax
```
scope scope-name {
    interface [ interface-names ];
    prefix destination-prefix;
}
```

Hierarchy Level
```
[edit logical-systems logical-system-name routing-instances routing-instance-name
    routing-options multicast],
[edit logical-systems logical-system-name routing-options multicast],
[edit routing-instances routing-instance-name routing-options multicast],
[edit routing-options multicast]
```

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure multicast scoping.

Options ***scope-name***—Name of the multicast scope.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Configuring Multicast Snooping*

Copyright © 2017, Juniper Networks, Inc.

scope-policy

Syntax scope-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-options multicast]

NOTE: You can configure a scope policy at these two hierarchy levels only. You cannot apply a scope policy to a specific routing instance, because all scoping policies are applied to all routing instances. However, you can apply the **scope** statement to a specific routing instance at the **[edit routing-instances *routing-instance-name* routing-options multicast]** or **[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast]** hierarchy level.

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.

Description Apply policies for scoping. The policy must be correctly configured at the **edit policy-options policy-statement** hierarchy level.

Options ***policy-names***—Name of one or more multicast scope policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- scope on page 157

Copyright © 2017, Juniper Networks, Inc.

source (Source-Specific Multicast)

Syntax source [*addresses*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast ssm-map *ssm-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast ssm-map *ssm-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast ssm-map *ssm-map-name*],
[edit routing-options multicast ssm-map *ssm-map-name*]

Release Information Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Specify IPv4 or IPv6 source addresses for an SSM map.

Options ***addresses***—IPv4 or IPv6 source addresses.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To view this statement in the configuration.

Related Documentation

- *Example: Configuring SSM Mapping*

Copyright © 2017, Juniper Networks, Inc.

source-routing

Syntax

```
source-routing {
  (ip | ipv6)
}
```

Hierarchy Level [edit routing-options]

Release Information

Statement for IPv6 introduced in Junos OS Release 8.2.
Statement for IPv4 introduced in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.

Description Enable source routing.

Source routing allows a sender of a packet to partially or completely specify the route the packet takes through the network. In contrast, in non-source routing protocols, routers in the network determine the path based on the packet's destination.

NOTE: We recommend that you not use source routing. Instead, we recommend that you use policy-based routing or filter-based forwarding to route packets based on source addresses.

Default Disabled

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Filter-Based Forwarding on the Source Address*

Copyright © 2017, Juniper Networks, Inc.

ssm-groups

Syntax ssm-groups [*ip-addresses*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure source-specific multicast (SSM) groups.

By default, the SSM group multicast address is limited to the IP address range from 232.0.0.0 through 232.255.255.255. However, you can extend SSM operations into another Class D range by including the **ssm-groups** statement in the configuration. The default SSM address range from 232.0.0.0 through 232.255.255.255 cannot be used in the **ssm-groups** statement. This statement is for adding other multicast addresses to the default SSM group addresses. This statement does not override the default SSM group address range.

IGMPv3 supports SSM groups. By utilizing inclusion lists, only sources that are specified send to the SSM group.

Options ***ip-addresses***—List of one or more additional SSM group addresses separated by a space.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Source-Specific Multicast Groups with Any-Source Override*

Copyright © 2017, Juniper Networks, Inc.

ssm-map (Routing Options Multicast)

Syntax

```
ssm-map ssm-map-name {
    policy [ policy-names ];
    source [ addresses ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced in Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure SSM mapping.

Options

ssm-map-name—Name of the SSM map.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring SSM Mapping*

Copyright © 2017, Juniper Networks, Inc.

static (Routes)

Syntax

```
static {
  defaults {
    static-options;
  }
  rib-group group-name;
  route destination-prefix {
    next-hop address;
    next-hop options;
    qualified-next-hop address {
      metric metric;
      preference preference;
    }
    static-options;
  }
}
```

Hierarchy Level [edit routing-options],
[edit routing-options rib *routing-table-name*]

Release Information Statement introduced in Junos OS Release 11.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure static routes to be installed in the routing table. You can specify any number of routes within a single **static** statement, and you can specify any number of **static** options in the configuration.

Options

defaults—Specify global static route options. These options only set default attributes inherited by all newly created static routes. These are treated as global defaults and apply to all the static routes you configure in the **static** statement. This part of the **static** statement is optional.

route *destination-prefix*—Destination of the static route.

- **defaults**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.
- ***destination-prefix/prefix-length*—*destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.
- **next-hop *address***—Reach the next-hop routing device by specifying an IP address, an interface name, or an ISO network entity title (NET).
- ***nsap-prefix*—*nsap-prefix*** is the network service access point (NSAP) address for ISO.

next-hop options—Additional information for how to manage forwarding of packets to the next hop.

- **discard**—Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

Copyright © 2017, Juniper Networks, Inc.

- ***iso-net***—Reach the next-hop routing device by specifying an ISO NSAP.
- **next-table *routing-table-name***—Name of the next routing table to the destination.
- **receive**—Install a receive route for this destination into the routing table.
- **reject**—Do not forward packets addressed to this destination. Instead, drop the packets, send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

static-options—(Optional under **route**) Additional information about static routes, which is included with the route when it is installed in the routing table.

You can specify one or more of the following in ***static-options***. Each of the options is explained separately.

- **(active | passive)**;
- **(install | no-install)**;
- **(metric | metric2 | metric3 | metric4)** ***value*** **<type** ***type*****>**;
- **(preference | preference2 | color | color2)** ***preference*** **<type** ***type*****>**;
- **(resolve | no-resolve)**;
- **(no-retain | retain)**;

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Static Routing on page 4

Copyright © 2017, Juniper Networks, Inc.

subscriber-leave-timer

Syntax subscriber-leave-timer *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast interface *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast interface *interface-name*],
[edit routing-options multicast interface *interface-name*]

Release Information
Statement introduced in Junos OS Release 9.2.
Statement introduced in Junos OS Release 9.2 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message.

Options
seconds—Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message. Specifying a value of 0 results in an immediate update. This is the same as if the statement were not configured.

Range: 0 through 30

Default: 0 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

tag (Routing Options)

Syntax tag *metric* type *number*;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (aggregate | generate | static) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options aggregate | generate | static) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)],
[edit routing-options (aggregate | generate | static) (defaults | route)],
[edit routing-options rib *routing-table-name* (aggregate | generate | static) (defaults | route)]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Associate a tag with a static, aggregate, or generated route.

Default No tag strings are associated with routes.

Options
metric—Tag metric.
Range: 0 through 4,294,967,295

type *number*—Tag type.
Range: 1 through 16

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Summarizing Static Routes Through Route Aggregation*
- aggregate on page 57
- generate on page 93
- *static*

Copyright © 2017, Juniper Networks, Inc.

threshold (Multicast Forwarding Cache)

Syntax

```
threshold {
  log-warning value;
  suppress value;
  reuse value;
  mvpn-rpt-suppress value;
  mvpn-rpt-reuse value;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast forwarding-cache],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast forwarding-cache family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-options multicast forwarding-cache],
[edit logical-systems *logical-system-name* routing-options multicast forwarding-cache family (inet | inet6)],
[edit routing-instances *routing-instance-name* routing-options multicast forwarding-cache],
[edit routing-instances *routing-instance-name* routing-options multicast forwarding-cache (inet | inet6)],
[edit routing-options multicast forwarding-cache],
[edit routing-options multicast forwarding-cache family (inet | inet6)]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.2 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure the suppression, reuse, and warning log message thresholds for multicast forwarding cache limits. You can configure the thresholds globally for the multicast forwarding cache or individually for the IPv4 and IPv6 multicast forwarding caches. Configuring the **threshold** statement globally for the multicast forwarding cache or including the **family** statement to configure the thresholds for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.

When general forwarding-cache suppression is active, the multicast forwarding-cache prevents forwarding traffic on the shared RP tree (RPT). At the same time, MVPN (*,G) forwarding states are not created for new RPT c-mcast entires, and , (*,G) installed by BGP-MVPN protocol are deleted. When general forwarding-cache suppression ends, BGP-MVPN (*,G) entries are re-added in the RIB and restored to the FIB (up to the MVPN (*,G) limit).

When MVPN RPT suppression is active, for all PE routers in excess of the threshold (including RP PEs), MVPN will not add new (*,G) forwarding entries to the forwarding-cache. Changes are visible once the entries in the current forwarding-cache have timed out or are deleted.

To use **mvpn-rpt-suppress** and/or **mvpn-rpt-reuse**, you must first configure the general **suppress** threshold. If **suppress** is configured but **mvpn-rpt-suppress** is not, both **mvpn-rpt-suppress** and **mvpn-rpt-reuse** will inherit *and use* the value set for the general **suppress**.

Copyright © 2017, Juniper Networks, Inc.

Options

reuse or **mvpn-rpt-reuse***value* (Optional) Value at which to begin creating new multicast forwarding cache entries. If configured, this number should be less than the **suppress** value.

Range: 1 through 200,000

suppress or **mvpn-rpt-suppress***value* —Value at which to begin suppressing new multicast forwarding cache entries. This value is mandatory. This number should be greater than the **reuse** value.

Range: 1 through 200,000

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Examples: Configuring the Multicast Forwarding Cache*
- *show multicast forwarding-cache statistics*

timeout (Flow Maps)

Syntax

timeout (never non-discard-entry-only | *minutes*);

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast flow-map *flow-map-name*],
[edit routing-options multicast flow-map *flow-map-name*]

Release Information

Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure the timeout value for multicast forwarding cache entries associated with the flow map.

Options

minutes—Length of time that the forwarding cache entry remains active.

Range: 1 through 720

never non-discard-entry-only—Specify that the forwarding cache entry always remain active. If you omit the **non-discard-entry-only** option, all multicast forwarding entries, including those in forwarding and pruned states, are kept forever. If you include the **non-discard-entry-only** option, entries with forwarding states are kept forever, and entries with pruned states time out.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

timeout (Multicast)

Syntax timeout *minutes* <family (inet | inet6)>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast forwarding-cache],
[edit logical-systems *logical-system-name* routing-options multicast forwarding-cache],
[edit routing-instances *routing-instance-name* routing-options multicast forwarding-cache],
[edit routing-options multicast forwarding-cache]

Release Information Statement introduced in Junos OS Release 8.2.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the timeout value for multicast forwarding cache entries.

Options ***minutes***—Length of time that the forwarding cache limit remains active.

Range: 1 through 720

family (inet | inet6)—(Optional) Apply the configured timeout to either IPv4 or IPv6 multicast forwarding cache entries. Configuring the **timeout** statement globally for the multicast forwarding cache or including the **family** statement to configure the timeout value for the IPv4 and IPv6 multicast forwarding caches are mutually exclusive.

Default: By default, the configured timeout applies to both IPv4 and IPv6 multicast forwarding cache entries.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring the Multicast Forwarding Cache*

Copyright © 2017, Juniper Networks, Inc.

traceoptions

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options],
[edit routing-options flow],
[edit routing-options multicast]

Release Information

Statement introduced before Junos OS Release 7.4.
nsr-synchronization flag for BGP, IS-IS, LDP, and OSPF added in Junos OS Release 8.4.
nsr-synchronization and **nsr-packet** flags for BFD sessions added in Junos OS Release 8.5.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
nsr-synchronization flag for RIP and RIPng added in Junos OS Release 9.0.
nsr-synchronization flag for Layer 2 VPNs and VPLS added in Junos OS Release 9.1.
nsr-synchronization flag for PIM added in Junos OS Release 9.3.
nsr-synchronization flag for MPLS added in Junos OS Release 10.1.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
nsr-synchronization flag for MSDP added in Junos OS Release 12.1.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Define tracing operations that track all routing protocol functionality in the routing device.

To specify more than one tracing operation, include multiple **flag** statements.

Default

If you do not include this statement, no global tracing operations are performed.

Options

Values:

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place global routing protocol tracing output in the file **routing-log**.

Copyright © 2017, Juniper Networks, Inc.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals
- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-packet**—Detailed trace information for BFD nonstop active routing only
- **nsr-synchronization**—Tracing operations for nonstop active routing
- **nsr-synchronization**—Nonstop active routing synchronization
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

no-world-readable—(Optional) Prevent any user from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

Copyright © 2017, Juniper Networks, Inc.

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- *Example: Tracing Global Routing Protocol Operations*

Copyright © 2017, Juniper Networks, Inc.

upstream-interface

Syntax upstream-interface [*interface-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast pim-to-igmp-proxy],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast pim-to-mld-proxy],
[edit logical-systems *logical-system-name* routing-options multicast pim-to-igmp-proxy],
[edit logical-systems *logical-system-name* routing-options multicast pim-to-mld-proxy],
[edit routing-instances *routing-instance-name* routing-options multicast pim-to-igmp-proxy],
[edit routing-instances *routing-instance-name* routing-options multicast pim-to-mld-proxy],
[edit routing-options multicast pim-to-igmp-proxy],
[edit routing-options multicast pim-to-mld-proxy]

Release Information
Statement introduced in Junos OS Release 9.6.
Statement introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 11.3 for the QFX Series.
Statement introduced in Junos OS Release 12.3 for ACX Series routers.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure at least one, but not more than two, upstream interfaces on the rendezvous point (RP) routing device that resides between a customer edge–facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain. The RP routing device translates PIM join or prune messages into corresponding IGMP report or leave messages (if you include the **pim-to-igmp-proxy** statement), or into corresponding MLD report or leave messages (if you include the **pim-to-mld-proxy** statement). The routing device then proxies the IGMP or MLD report or leave messages to one or both upstream interfaces to forward IPv4 multicast traffic (for IGMP) or IPv6 multicast traffic (for MLD) across the PIM domains.

Options
interface-names—Names of one or two upstream interfaces to which the RP routing device proxies IGMP or MLD report or leave messages for transmission of multicast traffic across PIM domains. You can specify a maximum of two upstream interfaces on the RP routing device. To configure a set of two upstream interfaces, specify the full interface names, including all physical and logical address components, within square brackets ([ ]).

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Configuring PIM-to-IGMP Message Translation*
- *Configuring PIM-to-MLD Message Translation*

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 3

Routing Options Operational Commands

- clear ipv6 neighbors
- show as-path
- show as-path domain
- show as-path summary
- show ipv6 neighbors
- show ipv6 router-advertisement
- show route
- show route active-path
- show route all
- show route aspath-regex
- show route best
- show route brief
- show route community
- show route community-name
- show route damping
- show route detail
- show route exact
- show route export
- show route extensive
- show route flow validation
- show route forwarding-table
- show route inactive-path
- show route inactive-prefix
- show route instance
- show route label
- show route label-switched-path
- show route martians
- show route next-hop

Copyright © 2017, Juniper Networks, Inc.

- show route no-community
- show route protocol
- show route range
- show route receive-protocol
- show route resolution
- show route snooping
- show route source-gateway
- show route summary
- show route table
- show route terse

Copyright © 2017, Juniper Networks, Inc.

clear ipv6 neighbors

Syntax
clear ipv6 neighbors
<all | host *hostname*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.3 for EX Series switches.
Command introduced in Junos OS Release 12.2 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Clear IPv6 neighbor cache information.

Options
none—Clear all IPv6 neighbor cache information.

all—(Optional) Clear all IPv6 neighbor cache information.

host *hostname*—(Optional) Clear the information for the specified IPv6 neighbors.

Required Privilege Level
view

Related Documentation
- show ipv6 neighbors on page 187

List of Sample Output
clear ipv6 neighbors on page 177

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
```

show as-path

List of Syntax
Syntax on page 178
Syntax (EX Series Switches) on page 178

Syntax
show as-path
<brief | detail>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show as-path
<brief | detail>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Display the distribution of autonomous system (AS) paths that the local routing device is using (usually through the routing table). Use this command to debug problems for AS paths and to understand how AS paths have been manipulated through a policy (through the **as-path-prepend** action) or through aggregation.

AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.

Options
none—Display basic information about AS paths that the local routing device is using (same as brief).

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show as-path summary on page 185

List of Sample Output
show as-path on page 179
show as-path detail on page 180

Output Fields
Table 12 on page 179 lists the output fields for the **show as-path** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 12: show as-path Output Fields

Field Name	Field Description	Level of Output
Total AS paths	Total number of AS paths.	**brief** none
Bucket	Bucket number.	All levels
Count	Number of AS path entries in this bucket.	All levels
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. • **Atomic**—Route is an aggregate of several route prefixes. • **Aggregator**—Routing device has summarized a range of prefixes.	All levels
domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.	**detail**
neighbor as	AS peer address.	**detail**
length	Length of the AS path.	**detail**
segments	Length of the AS segment descriptor.	**detail**
references	Path reference count.	**detail**

Sample Output

show as-path

```
user@host> show as-path
Total AS paths: 30382
  Bucket 0    Count: 36
    I
    14203 2914 174 31752 I
    14203 2914 701 21512 I
    14203 2914 1239 26632 I
    14203 2914 1239 29704 I
    14203 2914 4323 10248 I
    14203 2914 4766 23560 I
    14203 2914 6395 32776 I
    14203 2914 7911 11272 I
    14203 2914 12180 18440 I
    14203 2914 17408 17416 I
    14203 2914 701 702 24586 I
    14203 2914 1239 4657 9226 I
    14203 2914 1239 7132 16394 I
    14203 2914 1299 8308 34826 I
    14203 2914 3320 5603 28682 I
```

Copyright © 2017, Juniper Networks, Inc.

```
    14203 2914 3491 1680 33802 I
    14203 2914 3549 7908 27658 I
    14203 2914 3549 20804 30730 I
    14203 2914 7018 2687 9226 I
    14203 2914 174 9318 9318 23564 I
    14203 2914 701 3786 3786 23564 I
    14203 2914 701 4761 4795 9228 I
    14203 2914 1239 7132 5673 18444 I
    14203 2914 3491 20485 24588 24588 I
    14203 2914 5511 2200 1945 2060 I
    14203 2914 7911 14325 14325 14348 I
    14203 2914 701 4637 9230 9230 9230 I
    14203 2914 6395 14 14 14 14 I
    14203 2914 9299 6163 6163 6163 6163 9232 I
    14203 2914 3356 3356 3356 3356 3356 11955 21522 I
    14203 2914 9837 9837 9219 I Aggregator: 9219 202.27.91.253
    14203 2914 174 30209 30222 30222 30222 ?
    14203 2914 1299 5377 I (Atomic) Aggregator: 5377 193.219.192.22
    14203 2914 4323 36097 I (Atomic) Aggregator: 36097 216.69.252.254
    14203 2914 209 2516 17676 23813 I (Atomic) Aggregator: 23813 219.127.233.66
  Bucket 1    Count: 28
    14203 2914 35847 I
    14203 2914 174 19465 I
    14203 2914 174 35849 I
    14203 2914 2828 32777 I
    14203 2914 4323 14345 I
    14203 2914 4323 29705 I
    14203 2914 6395 32777 I

...
```

show as-path detail

```
user@host> show as-path detail
Total AS paths: 30410
  Bucket 0    Count: 36
    AS path: I
      domain 0, length 0, segments 0, references 54
    AS path: 14203 2914 174 31752 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 2
    AS path: 14203 2914 701 21512 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 2
    AS path: 14203 2914 1239 26632 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 2
    AS path: 14203 2914 1239 29704 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 2
    AS path: 14203 2914 4323 10248 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 2
    AS path: 14203 2914 4766 23560 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 2
    AS path: 14203 2914 6395 32776 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 3
    AS path: 14203 2914 7911 11272 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 2
    AS path: 14203 2914 12180 18440 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 3
    AS path: 14203 2914 17408 17416 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 3
    AS path: 14203 2914 701 702 24586 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 3
    AS path: 14203 2914 1239 4657 9226 I
```


Copyright © 2017, Juniper Networks, Inc.

```
      domain 1, neighbor as: 14203, length 5, segments 1, references 7
    AS path: 14203 2914 1239 7132 16394 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 2
    AS path: 14203 2914 1299 8308 34826 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 2
    AS path: 14203 2914 3320 5603 28682 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 2
    AS path: 14203 2914 3491 1680 33802 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 2
    AS path: 14203 2914 3549 7908 27658 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 2
    AS path: 14203 2914 3549 20804 30730 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 2
    AS path: 14203 2914 7018 2687 9226 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 3
    AS path: 14203 2914 174 9318 9318 23564 I
      domain 1, neighbor as: 14203, length 6, segments 1, references 2
    AS path: 14203 2914 701 3786 3786 23564 I
      domain 1, neighbor as: 14203, length 6, segments 1, references 2
    AS path: 14203 2914 701 4761 4795 9228 I
      domain 1, neighbor as: 14203, length 6, segments 1, references 14
    AS path: 14203 2914 1239 7132 5673 18444 I
      domain 1, neighbor as: 14203, length 6, segments 1, references 2
    AS path: 14203 2914 3491 20485 24588 24588 I
      domain 1, neighbor as: 14203, length 6, segments 1, references 4
    AS path: 14203 2914 5511 2200 1945 2060 I
      domain 1, neighbor as: 14203, length 6, segments 1, references 2
    AS path: 14203 2914 7911 14325 14325 14348 I
      domain 1, neighbor as: 14203, length 6, segments 1, references 2
    AS path: 14203 2914 701 4637 9230 9230 9230 I
      domain 1, neighbor as: 14203, length 7, segments 1, references 3
    AS path: 14203 2914 6395 14 14 14 14 I
      domain 1, neighbor as: 14203, length 7, segments 1, references 10
...
```

show as-path domain

List of Syntax
Syntax on page 182
Syntax (EX Series Switches) on page 182

Syntax
show as-path domain
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show as-path domain

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display autonomous system (AS) path domain information.

Options
none—(Optional) Display AS path domain information for all routing instances.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show as-path domain on page 184

Output Fields
Table 13 on page 182 lists the output fields for the **show as-path domain** command. Output fields are listed in the approximate order in which they appear

Table 13: show as-path domain Output Fields

Field Name	Field Description
Domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.
Primary	Primary AS number.
References	Path reference count.
***Number* Paths**	Number of known AS paths.
Flags	Information about the AS path: • **ASLoop**—Path contains an AS loop. • **Atomic**—Path includes the ATOMIC_AGGREGATE path attribute. • **Local**—Path was created by local aggregation. • **Master**—Path was created by the master routing instance.
Local AS	AS number of the local routing device.

Copyright © 2017, Juniper Networks, Inc.

Table 13: show as-path domain Output Fields *(continued)*

Field Name	Field Description
Loops	How many times this AS number can appear in an AS path.

Copyright © 2017, Juniper Networks, Inc.

Sample Output

show as-path domain

```
user@host> show as-path domain
Domain: 1        Primary: 10458
    References:          3 Paths:      30383
    Flags: Master
    Local AS: 10458  Loops: 1
```


Copyright © 2017, Juniper Networks, Inc.

show as-path summary

List of Syntax
Syntax on page 185
Syntax (EX Series Switches) on page 185

Syntax
show as-path summary
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show as-path summary

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display autonomous system (AS) path summary information.

AS paths are stored in a hash table. A hash table is one method for fast lookup. Each entry in the table is called a bucket. Junos OS computes a hash value that indicates in which bucket the AS path is stored. The AS paths are dispersed among the hash buckets so that a manageable number of AS paths is stored in each bucket. Only unique AS paths are stored. Duplicate AS paths increase a reference count, but do not increase the number of AS paths stored in the hash table.

Options
none—(Optional) Display AS path summary information for all routing instances.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show as-path on page 178

List of Sample Output
show as-path summary on page 186

Output Fields
Table 14 on page 185 lists the output fields for the **show as-path summary** command. Output fields are listed in the approximate order in which they appear.

Table 14: show as-path summary Output Fields

Field Name	Field Description
AS Paths	Number of AS paths.
Buckets	Number of hash buckets in use.
Max	Maximum number of AS path entries per bucket.
Min	Minimum number of AS path entries per bucket.
Avg	Average number of AS path entries per bucket.

Copyright © 2017, Juniper Networks, Inc.

Table 14: show as-path summary Output Fields *(continued)*

Field Name	Field Description
Std deviation	Standard deviation of AS path entries per bucket.

Sample Output

show as-path summary

```
user@host> show as-path summary
AS Paths  Buckets      Max     Min     Avg   Std deviation
   30425     1024       95      12      29        6.481419
```


Copyright © 2017, Juniper Networks, Inc.

show ipv6 neighbors

Syntax show ipv6 neighbors

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.3 for EX Series switches.
Command introduced in Junos OS Release 12.2 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display information about the IPv6 neighbor cache.

NOTE: Starting with Junos OS Release 16.1, **show ipv6 neighbors** command does not display the underlying ifl information if **enhanced-convergence** statement at [**edit irb unit** *unit-number*] hierarchy level and **enhanced-ip** statement at [**edit chassis network-services**] hierarchy level is configured for the destination interface IRB.

Options This command has no options.

Required Privilege Level view

Related Documentation
- clear ipv6 neighbors on page 177

List of Sample Output
show ipv6 neighbors on page 188
show ipv6 neighbors on page 188

Output Fields Table 15 on page 187 describes the output fields for the **show ipv6 neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 15: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: **up**, **down**, **incomplete**, **reachable**, **stale**, or **unreachable**.
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: **yes** or **no**.
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: **yes** or **no**.

Copyright © 2017, Juniper Networks, Inc.

Table 15: show ipv6 neighbors Output Fields *(continued)*

Field Name	Field Description
Interface	Name of the interface.

Sample Output

show ipv6 neighbors

```
user@host> show ipv6 neighbors
IPv6 Address                 Linklayer Address  State       Exp Rtr Secure
Interface
2001:db8:0:1:2a0:a514:0:24c  00:05:85:8f:c8:bd  stale       546 yes no
fe-1/2/0.1
fe80::2a0:a514:0:24c         00:05:85:8f:c8:bd  stale       258 yes no
fe-1/2/0.1
fe80::2a0:a514:0:64c         00:05:85:8f:c8:bd  stale       111 yes no
fe-1/2/1.5
fe80::2a0:a514:0:a4c         00:05:85:8f:c8:bd  stale       327 yes no
fe-1/2/2.9
```

show ipv6 neighbors

The command displaying the underlying l2 ifl information when **enhanced-convergence** statement and **enhanced-ip** statement is not configured.

```
IPv6 Address                 Linklayer Address  State       Exp Rtr Secure
Interface
23::23:0:0:2                 00:00:23:00:00:02  reachable   0   no  no      irb.0
 [xe-2/2/0.0]
```

The command not displaying the underlying l2 ifl information when **enhanced-convergence** statement and **enhanced-ip** statement is configured.

```
IPv6 Address                 Linklayer Address  State       Exp Rtr Secure
Interface
23::23:0:0:2                 00:00:23:00:00:02  reachable   0   no  no      irb.0
```


Copyright © 2017, Juniper Networks, Inc.

show ipv6 router-advertisement

Syntax

show ipv6 router-advertisement
<conflicts>
<interface *interface*>
<logical-system (all | *logical-system-name*)>
<prefix *prefix/prefix length*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 12.2 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display information about IPv6 router advertisements, including statistics about messages sent and received on interfaces, and information received from advertisements from other routers.

Options

none—Display all IPv6 router advertisement information for all interfaces.

conflicts—(Optional) Display only the IPv6 router advertisement information that is conflicting.

interface *interface*—(Optional) Display IPv6 router advertisement information for the specified interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix *prefix/prefix length*—(Optional) Display IPv6 router advertisement information for the specified prefix.

Additional Information

The display identifies conflicting information by enclosing the value the router is advertising in brackets.

Required Privilege Level

view

Related Documentation

- *clear ipv6 router-advertisement*

List of Sample Output

show ipv6 router-advertisement on page 190
show ipv6 router-advertisement conflicts on page 191
show ipv6 router-advertisement prefix on page 191

Output Fields

Table 16 on page 189 describes the output fields for the **show ipv6 router-advertisement** command. Output fields are listed in the approximate order in which they appear.

Table 16: show ipv6 router-advertisement Output Fields

Field Name	Field Description
Interface	Name of the interface.

Copyright © 2017, Juniper Networks, Inc.

Table 16: show ipv6 router-advertisement Output Fields *(continued)*

Field Name	Field Description
Advertisements sent	Number of router advertisements sent and the elapsed time since they were sent.
Solicits received	Number of solicitation messages received.
Advertisements received	Number of router advertisements received.
Advertisements from	Names of interfaces from which router advertisements have been received and the elapsed time since the last one was received.
Managed	Managed address configuration flag: **0** (stateless) or **1** (stateful).
Other configuration	Other stateful configuration flag: **0** (stateless) or **1** (stateful).
Reachable time	Time that a node identifies a neighbor as reachable after receiving a reachability confirmation, in milliseconds.
Default lifetime	Default lifetime, in seconds: from 0 seconds to 18.2 hours. A setting of 0 indicates that the router is not a default router.
Retransmit timer	Time between retransmitted Neighbor Solicitation messages, in milliseconds.
Current hop limit	Configured current hop limit.
Prefix	Name and length of the prefix.
Valid lifetime	How long the prefix remains valid for onlink determination.
Preferred lifetime	How long the prefix generated by stateless autoconfiguration remains preferred.
On link	Onlink flag: **0** (not onlink) or **1** (onlink).
Autonomous	Autonomous address configuration flag: **0** (not autonomous) or **1** (autonomous).

Sample Output

show ipv6 router-advertisement

```
user@host> show ipv6 router-advertisement
Interface: fe-0/1/1.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 0
Interface: fxp0.0
  Advertisements sent: 0
  Solicits received: 0
  Advertisements received: 1
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:00:13 ago
    Managed: 0
```


Copyright © 2017, Juniper Networks, Inc.

```
    Other configuration: 0 [1]
      Reachable time: 0 ms
      Default lifetime: 1800 sec
      Retransmit timer: 0 ms
      Current hop limit: 64
```

show ipv6 router-advertisement conflicts

```
user@host> show ipv6 router-advertisement conflicts
Interface: fxp0.0
  Advertisement from fe80::2d0:b7ff:fe1e:7b0e, heard 00:01:08 ago
    Other configuration: 0 [1]
```

show ipv6 router-advertisement prefix

```
user@host> show ipv6 router-advertisement prefix 2001:db8:8040::/16
Interface: fe-0/1/3.0
  Advertisements sent: 3, last sent 00:04:11 ago
  Solicits received: 0
  Advertisements received: 3
  Advertisement from fe80::290:69ff:fe9a:5403, heard 00:00:05 ago
    Managed: 0
    Other configuration: 0
    Reachable time: 0 ms
    Default lifetime: 180 sec [1800 sec]
    Retransmit timer: 0 ms
    Current hop limit: 64
    Prefix: 2001:db8:8040:1::/64
      Valid lifetime: 2592000 sec
      Preferred lifetime: 604800 sec
      On link: 1
      Autonomous: 1
```

Copyright © 2017, Juniper Networks, Inc.

show route

List of Syntax
Syntax on page 192
Syntax (EX Series Switches) on page 192

Syntax

show route
<all>
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>
<private>

Syntax (EX Series Switches)

show route
<all>
<*destination-prefix*>
<private>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Option **private** introduced in Junos OS Release 9.5.
Option **private** introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.
Option **display-client-data** introduced in Junos OS Release 16.2R1 on MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series routers.

Description

Display the active entries in the routing tables.

Options

none—Display brief information about all active entries in the routing tables.

all—(Optional) Display information about all routing tables, including private, or internal, routing tables.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

private—(Optional) Display information only about all private, or internal, routing tables.

display-client-data—(Optional) Display client id and cookie information for routes installed by rpd client applications.

Required Privilege Level

view

Related Documentation

- *Understanding IS-IS Configuration*
- *Example: Configuring IS-IS*
- *Examples: Configuring Internal BGP Peering*
- *Examples: Configuring External BGP Peering*

Copyright © 2017, Juniper Networks, Inc.

- *Examples: Configuring OSPF Routing Policy*
- *Verifying and Managing Junos OS Enhanced Subscriber Management*

List of Sample Output

show route on page 196
show route (VPN) on page 197
show route (with Destination Prefix) on page 197
show route destination-prefix detail on page 197
show route extensive on page 197
show route extensive (ECMP) on page 198
show route (Enhanced Subscriber Management) on page 198
show route (IPv6 Flow Specification) on page 198
show route display-client-data detail on page 199

Output Fields

Table 17 on page 193 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

Table 17: show route Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active). • **holddown** (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly. However, if you have configured advertisement of multiple routes (with the **add-path** or **advertise-inactive** statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong. • **hidden** (routes that are not used because of a routing policy).

Copyright © 2017, Juniper Networks, Inc.

Table 17: show route Output Fields *(continued)*

Field Name	Field Description

destination-prefix	Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as:
	• ***MPLS-label*** (for example, 80001).
	• ***interface-name*** (for example, ge-1/0/2).
	• ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96):
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.

[***protocol*, *preference*** **]**	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
weeks:days hours:minutes:seconds	How long the route been known (for example, **2w4d 13:11:14**, or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).
metric	Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
localpref	Local preference value included in the route.
from	Interface from which the route was received.

Copyright © 2017, Juniper Networks, Inc.

Table 17: show route Output Fields *(continued)*

Field Name	Field Description
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.
Route Labels	Stack of labels carried in the BGP route update.
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the destination is **Discard**, traffic is dropped.

Copyright © 2017, Juniper Networks, Inc.

Table 17: show route Output Fields *(continued)*

Field Name	Field Description
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • **Weight**—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • **Balance**—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. • ***lsp-path-name***—Name of the LSP used to reach the next hop. • ***label-action***—MPLS label and operation occurring at the next hop. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label). For VPNs, expect to see multiple **push** operations, corresponding to the inner and outer labels required for VPN routes (in the case of a direct PE-to-PE connection, the VPN route would have the inner label push only).
Private unicast	(Enhanced subscriber management for MX Series routers) Indicates that an access-internal route is managed by enhanced subscriber management. By contrast, access-internal routes *not* managed by enhanced subscriber management are displayed with associated next-hop and media access control (MAC) address information.
balance	Distribution of the load based on the underlying operational interface bandwidth for equal-cost multipaths (ECMP) across the nexthop gateways in percentages.

Sample Output

show route

```
user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
                   *[MVPN/70] 19:53:41, metric2 1
                      Indirect
1:65500:1:10.0.0.40/240
                   *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
                      AS path: I
                    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
                    [BGP/170] 19:53:26, localpref 100, from 10.0.0.33
                      AS path: I
                    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
                   *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
                      AS path: I
                    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
                    [BGP/170] 19:53:25, localpref 100, from 10.0.0.33
                      AS path: I
                    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
```


Copyright © 2017, Juniper Networks, Inc.

show route (VPN)

The following sample output shows a VPN route with composite next hops enabled. The first **Push** operation corresponds to the outer label. The second **Push** operation corresponds to the inner label.

```
user@host> show route 192.0.2.0

13979:665001.inet.0: 871 destinations, 3556 routes (871 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.0/24        @[BGP/170] 00:28:32, localpref 100, from 10.9.9.160
                      AS path: 13980 ?, validation-state: unverified
                    > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                    [BGP/170] 00:28:28, localpref 100, from 10.9.9.169
                      AS path: 13980 ?, validation-state: unverified
                    > to 10.100.0.42 via ae2.0, Push 126016, Push 300368(top)
                    #[Multipath/255] 00:28:28, metric2 102
                    > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                      to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
```

show route (with Destination Prefix)

```
user@host> show route 172.16.0.0/12

inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.0/12       *[Static/5] 2w4d 12:54:27
                     > to 192.168.167.254 via fxp0.0
```

show route destination-prefix detail

```
user@host> show route 198.51.100.0 detail

inet.0: 15 destinations, 20 routes (15 active, 0 holddown, 0 hidden)
198.51.100.0/24 (2 entries, 2 announced)
        *BGP    Preference: 170/-101
         ...
         BGP-Static Preference: 4294967292
               Next hop type: Discard
                 Address: 0x9041ae4
                 Next-hop reference count: 2
                 State: <NoReadvrt Int Ext AlwaysFlash>
        Inactive reason: Route Preference
        Local AS:   200
        Age: 4d 1:40:40
        Validation State: unverified
        Task: RT
        Announcement bits (1): 2-BGP_RT_Background
        AS path: 4 5 6 I
```

show route extensive

```
user@host> show route extensive
v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
    *BGP   Preference: 170/-101
           PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group
```

Copyright © 2017, Juniper Networks, Inc.

```
203.0.113.1
           Next hop type: Indirect
           Address: 0x92455b8
           Next-hop reference count: 2
           Source: 10.0.0.30
           Protocol next hop: 10.0.0.40
           Indirect next hop: 2 no-forward
           State: <Active Int Ext>
                    Local AS: 64510 Peer AS: 64511
           Age: 3  Metric2: 1
          Validation State: unverified
           Task: BGP_64510.10.0.0.30+179
           Announcement bits (2): 0-PIM.v1 1-mvpn global task
           AS path: I (Originator) Cluster list:  10.0.0.30
           AS path:  Originator ID: 10.0.0.40
           Communities: target:64502:100
           Import Accepted
           Localpref: 100
           Router ID: 10.0.0.30
           Primary Routing Table bgp.mvpn.0
           Indirect next hops: 1
                Protocol next hop: 10.0.0.40 Metric: 1
                Indirect next hop: 2 no-forward
                Indirect path forwarding next hops: 1
                      Next hop type: Router
                      Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
                10.0.0.40/32 Originating RIB: inet.3
                      Metric: 1                       Node path count: 1
                      Forwarding nexthops: 1
                         Nexthop: 10.0.24.4 via lt-0/3/0.24
```

show route extensive (ECMP)

```
user@host> show route extensive
*IS-IS Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 1048577
                Address: 0xXXXXXXXXXX
                Next-hop reference count: YY
                Next hop: 172.16.50.2 via ae1.0 balance 43%, selected
                Session Id: 0x141
                Next hop: 192.0.2.2 via ae0.0 balance 57%
```

show route (Enhanced Subscriber Management)

```
user@host> show route
inet.0: 41 destinations, 41 routes (40 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

198.51.100.11/24    *[Access-internal/12] 00:00:08
                > to #0 10.0.0.1.93.65 via demux0.1073741824
198.51.100.12/24    *[Access-internal/12] 00:00:08
                Private unicast
```

show route (IPv6 Flow Specification)

```
user@host> show route
inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::10:255:185:19/128
                   *[Direct/0] 05:11:27
```


Copyright © 2017, Juniper Networks, Inc.

```
                    > via lo0.0
2001:db8::11:11:11:0/120
                   *[BGP/170] 00:28:58, localpref 100
                      AS path: 2000 I, validation-state: unverified
                    > to 2001:db8::13:14:2:2 via ge-1/1/4.0
2001:db8::13:14:2:0/120*[Direct/0] 00:45:07
                    > via ge-1/1/4.0
2001:db8::13:14:2:1/128*[Local/0] 00:45:18
                      Local via ge-1/1/4.0
fe80::2a0:a50f:fc71:71d5/128
                   *[Direct/0] 05:11:27
                    > via lo0.0
fe80::5e5e:abff:feb0:933e/128
                   *[Local/0] 00:45:18
                      Local via ge-1/1/4.0

inet6flow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::11:11:11:10/128,*,proto=6,dstport=80,srcport=65535/term:1
                   *[BGP/170] 00:28:58, localpref 100, from 2001:db8::13:14:2:2
                      AS path: 2000 I, validation-state: unverified
                      Fictitious
2001:db8::11:11:11:30/128,*,icmp6-type=128,len=100,dscp=10/term:2
                   *[BGP/170] 00:20:54, localpref 100, from 2001:db8::13:14:2:2
                      AS path: 2000 I, validation-state: unverified
                      Fictitious
```

show route display-client-data detail

```
user@host> show route 198.51.100.0/24 display-client-data detail
inet.0: 59 destinations, 70 routes (59 active, 0 holddown, 0 hidden)
198.51.100.0/24 (1 entry, 1 announced)
        State: <FlashAll>
        *BGP-Static Preference: 5/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xa5c2af8
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 1641
                Next hop: 192.0.2.1 via ge-2/1/1.0, selected
                Session Id: 0x160
                Protocol next hop: 192.0.2.1
                Indirect next hop: 0xa732cb0 1048621 INH Session ID: 0x17e
                State: <Active Int Ext AlwaysFlash NSR-incapable Programmed>
                Age: 3:13       Metric2: 0
                Validation State: unverified
                Announcement bits (3): 0-KRT 5-LDP 6-Resolve tree 3
                AS path: I
                Client id: 1, Cookie: 1
```

Copyright © 2017, Juniper Networks, Inc.

show route active-path

List of Syntax
Syntax on page 200
Syntax (EX Series Switches) on page 200

Syntax
show route active-path
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route active-path
<brief | detail | extensive | terse>

Release Information
Command introduced in Junos OS Release 8.0.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed.

Options
none—Display all active routes.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route active-path on page 200
show route active-path brief on page 201
show route active-path detail on page 201
show route active-path extensive on page 202
show route active-path terse on page 204

Output Fields
For information about output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route active-path

```
user@host> show route active-path

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32    *[Direct/0] 21:33:52
                    > via lo0.0
10.255.71.50/32    *[IS-IS/15] 00:18:13, metric 10
                    > to 172.16.100.1 via so-2/1/3.0
172.16.100.1/24      *[Direct/0] 00:18:36
                    > via so-2/1/3.0
```


Copyright © 2017, Juniper Networks, Inc.

```
172.16.100.1/32       *[Local/0] 00:18:41
                       Local via so-2/1/3.0
192.168.64.0/21    *[Direct/0] 21:33:52
                    > via fxp0.0
192.168.70.19/32   *[Local/0] 21:33:52
                       Local via fxp0.0
```

show route active-path brief

The output for the **show route active-path brief** command is identical to that for the **show route active-path** command. For sample output, see show route active-path on page 200.

show route active-path detail

```
user@host> show route active-path detail

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via lo0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

10.255.71.50/32 (1 entry, 1 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 397
                Next-hop reference count: 4
                Next hop: 172.16.100.1 via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:31      Metric: 10
                Task: IS-IS
                Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
                AS path: I

172.16.100.0/24 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:54
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

172.16.100.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Next-hop reference count: 11
                    Interface: so-2/1/3.0
                    State: ‹Active NoReadvrt Int›
                    Local AS:   200
                    Age: 21:59
                    Task: IF
                    Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                    AS path: I

192.168.64.0/21 (1 entry, 1 announced)
        *Direct Preference: 0
                    Next hop type: Interface
                    Next-hop reference count: 3
                    Next hop: via fxp0.0, selected
                    State: ‹Active Int›
                    Local AS:   200
                    Age: 21:37:10
                    Task: IF
                    Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                    AS path: I

192.168.70.19/32 (1 entry, 1 announced)
        *Local  Preference: 0
                    Next hop type: Local
                    Next-hop reference count: 11
                    Interface: fxp0.0
                    State: ‹Active NoReadvrt Int›
                    Local AS:   200
                    Age: 21:37:10
                    Task: IF
                    Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                    AS path: I
```

show route active-path extensive

```
user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
        *Direct Preference: 0
                    Next hop type: Interface
                    Next-hop reference count: 3
                    Next hop: via lo0.0, selected
                    State: ‹Active Int›
                    Local AS:   200
                    Age: 21:39:47
                    Task: IF
                    Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                    AS path: I

10.255.71.50/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.255.71.50/32 -> {172.16.100.1}
IS-IS level 2, LSP fragment 0
        *IS-IS  Preference: 15
                    Level: 1
                    Next hop type: Router, Next hop index: 397
```

 Copyright © 2017, Juniper Networks, Inc.

```
                Next-hop reference count: 4
                Next hop: 172.16.100.1 via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 24:08      Metric: 10
                Task: IS-IS
                Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
                AS path: I

172.16.100.1/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 24:31
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

172.16.100.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: so-2/1/3.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 24:36
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.64.0/21 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via fxp0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.70.19/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: fxp0.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
```

Copyright © 2017, Juniper Networks, Inc.

```
                        AS path: I
```

show route active-path terse

```
user@host> show route active-path terse

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.19/32    D   0                     >lo0.0
* 10.255.71.50/32    I  15         10          >172.16.100.1.
* 172.16.100.0/24       D   0                     >so-2/1/3.0
* 172.16.100.2/32       L   0                      Local
* 192.168.64.0/21    D   0                     >fxp0.0
* 192.168.70.19/32   L   0                      Local
```


Copyright © 2017, Juniper Networks, Inc.

show route all

List of Syntax
Syntax on page 205
Syntax (EX Series Switches) on page 205

Syntax
show route all
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route all

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display information about all routes in all routing tables, including private, or internal, tables.

Options
none—Display information about all routes in all routing tables, including private, or internal, tables.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show route brief on page 212
- show route detail on page 223

List of Sample Output
show route all on page 205

Output Fields
In Junos OS Release 9.5 and later, only the output fields for the **show route all** command display all routing tables, including private, or hidden, routing tables. The output field table of the show route command does not display entries for private, or hidden, routing tables in Junos OS Release 9.5 and later.

Sample Output

show route all

The following example displays a snippet of output from the **show route** command and then displays the same snippet of output from the **show route all** command:

```
user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0                  *[MPLS/0] 2d 02:24:39, metric 1
                      Receive
1                  *[MPLS/0] 2d 02:24:39, metric 1
                      Receive
2                  *[MPLS/0] 2d 02:24:39, metric 1
                      Receive
```

Copyright © 2017, Juniper Networks, Inc.

```
800017             *[VPLS/7] 1d 14:00:16
                    > via vt-3/2/0.32769, Pop
800018             *[VPLS/7] 1d 14:00:26
                    > via vt-3/2/0.32772, Pop

user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0                  *[MPLS/0] 2d 02:19:12, metric 1
                      Receive
1                  *[MPLS/0] 2d 02:19:12, metric 1
                      Receive
2                  *[MPLS/0] 2d 02:19:12, metric 1
                      Receive
800017             *[VPLS/7] 1d 13:54:49
                    > via vt-3/2/0.32769, Pop
800018             *[VPLS/7] 1d 13:54:59
                    > via vt-3/2/0.32772, Pop
vt-3/2/0.32769      [VPLS/7] 1d 13:54:49
                      Unusable
vt-3/2/0.32772      [VPLS/7] 1d 13:54:59
                      Unusable
```


Copyright © 2017, Juniper Networks, Inc.

show route aspath-regex

List of Syntax
Syntax on page 207
Syntax (EX Series Switches) on page 207

Syntax
show route aspath-regex *regular-expression*
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route aspath-regex *regular-expression*

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the entries in the routing table that match the specified autonomous system (AS) path regular expression.

Options
regular-expression—Regular expression that matches an entire AS path.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information
You can specify a regular expression as:

- An individual AS number
- A period wildcard used in place of an AS number
- An AS path regular expression that is enclosed in parentheses

You also can include the operators described in the table of AS path regular expression operators in the *Junos Policy Framework Configuration Guide*. The following list summarizes these operators:

- ***{m,n}***—At least *m* and at most *n* repetitions of the AS path term.
- ***{m}***—Exactly *m* repetitions of the AS path term.
- ***{m,}***—*m* or more repetitions of the AS path term.
- *******—Zero or more repetitions of an AS path term.
- ***+***—One or more repetitions of an AS path term.
- ***?***—Zero or one repetition of an AS path term.
- ***aspath_term | aspath_term***—Match one of the two AS path terms.

  When you specify more than one AS number or path term, or when you include an operator in the regular expression, enclose the entire regular expression in quotation marks. For example, to match any path that contains AS number 234, specify the following command:

  show route aspath-regex ".* 234 .*"

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level view

Related Documentation
- *Example: Using AS Path Regular Expressions*

List of Sample Output
show route aspath-regex (Matching a Specific AS Number) on page 208
show route aspath-regex (Matching Any Path with Two AS Numbers) on page 208

Output Fields For information about output fields, see the output field table for the show route command.

Sample Output

show route aspath-regex (Matching a Specific AS Number)

```
user@host> show route aspath-regex 65477
inet.0: 46411 destinations, 46411 routes (46409 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

111.222.1.0/25     *[BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65548 65536}) IGP
                     to 111.222.18.225 via fpa0.0(111.222.18.233)
111.222.1.128/25   *[IS-IS/15] 09:15:37, metric 37, tag 1
                     to 111.222.18.225 via fpa0.0(111.222.18.233)
                    [BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65548 65536}) IGP
                     to 111.222.18.225 via fpa0.0(111.222.18.233)
...
```

show route aspath-regex (Matching Any Path with Two AS Numbers)

```
user@host> show route aspath-regex ".* 234 3561 .*"

inet.0: 46351 destinations, 46351 routes (46349 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

9.20.0.0/17        *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 2685 2686 Incomplete
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
12.10.231.0/24     *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 5696 7369 IGP
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.64.32.0/19      *[BGP/170] 01:34:59, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 6327 IGP
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
...
```


Copyright © 2017, Juniper Networks, Inc.

show route best

List of Syntax
Syntax on page 209
Syntax (EX Series Switches) on page 209

Syntax
show route best *destination-prefix*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route best *destination-prefix*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

destination-prefix—Address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show route brief on page 212
- show route detail on page 223

List of Sample Output
show route best on page 209
show route best detail on page 210
show route best extensive on page 211
show route best terse on page 211

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route best

```
user@host> show route best 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32   *[OSPF/10] 1d 13:19:20, metric 2
                    > to 10.31.1.6 via ge-3/1/0.0
                      via so-0/3/0.0
```

Copyright © 2017, Juniper Networks, Inc.

```
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32   *[RSVP/7] 1d 13:20:13, metric 2
                    > via so-0/3/0.0, label-switched-path green-r1-r3

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.0/8         *[Direct/0] 2d 01:43:34
                    > via fxp2.0
                    [Direct/0] 2d 01:43:34
                    > via fxp1.0
```

show route best detail

```
user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                Next hop: via so-0/3/0.0
                State: <Active Int>
                Local AS:    69
                Age: 1d 13:20:06        Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 5
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100016
                State: <Active Int>
                Local AS:    69
                Age: 1d 13:20:59        Metric: 2
                Task: RSVP
                Announcement bits (1): 1-Resolve tree 2
                AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp2.0, selected
                State: <Active Int>
                Age: 2d 1:44:20
                Task: IF
                AS path: I
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Next hop: via fxp1.0, selected
                    State: <NotBest Int>
                    Inactive reason: No difference
                    Age: 2d 1:44:20
                    Task: IF
                    AS path: I
```

show route best extensive

The output for the **show route best extensive** command is identical to that for the **show route best detail** command. For sample output, see show route best detail on page 210.

show route best terse

```
user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.103/32   O  10          2            >10.31.1.6
                                                   so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.103/32   R   7          2            >so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.0.0.0/8         D   0                        >fxp2.0
                     D   0                        >fxp1.0
```

Copyright © 2017, Juniper Networks, Inc.

show route brief

List of Syntax
Syntax on page 212
Syntax (EX Series Switches) on page 212

Syntax
show route brief
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route brief
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display brief information about the active entries in the routing tables.

Options
none—Display all active entries in the routing table.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show route all on page 205
- show route best on page 209

List of Sample Output
show route brief on page 212

Output Fields
For information about output fields, see the Output Field table of the show route command.

Sample Output

show route brief

```
user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                      Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                    > via lo0.0
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.0.0/18      *[Static/5] 1w5d 20:30:29
                    > to 192.168.167.254 via fxp0.0
192.168.40.0/22    *[Static/5] 2w4d 13:11:14
```


Copyright © 2017, Juniper Networks, Inc.

```
                    > to 192.168.167.254 via fxp0.0
192.168.64.0/18     *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.164.0/22    *[Direct/0] 2w4d 13:11:14
                    > via fxp0.0
192.168.164.51/32   *[Local/0] 2w4d 13:11:14
                      Local via fxp0.0
207.17.136.192/32   *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16      *[Direct/0] 1w5d 20:30:28
                    > via fe-0/0/3.0
100.101.2.3/32      *[Local/0] 1w5d 20:30:28
                      Local via fe-0/0/3.0
172.16.233.5/32       *[OSPF/10] 1w5d 20:30:29, metric 1
                      MultiRecv
```

Copyright © 2017, Juniper Networks, Inc.

show route community

List of Syntax
Syntax on page 214
Syntax (EX Series Switches) on page 214

Syntax
show route community *as-number:community-value*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route community *as-number:community-value*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community.

Options
as-number:community-value—One or more community identifiers. ***as-number*** is the AS number, and ***community-value*** is the community identifier. When you specify more than one community identifier, enclose the identifiers in double quotation marks. Community identifiers can include wildcards.

For example:

```
user@host> show route table inet.0 protocol bgp community "12083:6015"
  community "12083:65551"
```

or

```
user@host> show route table inet.0 protocol bgp community [12083:6014
  12083:65551]
```

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information
Specifying the community option displays all routes matching the community found within the routing table. The community option does not limit the output to only the routes being advertised to the neighbor after any egress routing policy.

Required Privilege Level
view

Related Documentation
- show route detail on page 223

List of Sample Output
show route community on page 215

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Copyright © 2017, Juniper Networks, Inc.

Sample Output

show route community

```
user@host> show route community 234:80
inet.0: 46511 destinations, 46511 routes (46509 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.4.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                    AS Path: {666} 234 2548 1 IGP
                      to 192.156.169.1 via 192.156.169.14(so-0/0/0)
172.16.6.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                    AS Path: {666} 234 2548 568 721 Incomplete
                      to 192.156.169.1 via 192.156.169.14(so-0/0/0)
172.16.92.0/16         *[BGP/170] 03:33:06, localpref 100, from 131.103.20.49
                    AS Path: {666} 234 2548 1673 1675 1747 IGP
                      to 192.156.169.1 via 192.156.169.14(so-0/0/0)
```

Copyright © 2017, Juniper Networks, Inc.

show route community-name

List of Syntax Syntax on page 216
Syntax (EX Series Switches) on page 216

Syntax

```
show route community-name community-name
<brief | detail | extensive | terse>
<logical-system (all | logical-system-name)>
```

Syntax (EX Series Switches)

```
show route community-name community-name
<brief | detail | extensive | terse>
```

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community, specified by a community name.

Options

community-name—Name of the community.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level view

List of Sample Output show route community-name on page 216

Output Fields For information about output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route community-name

```
user@host> show route community-name red-com
inet.0: 17 destinations, 17 routes (16 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

instance1.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.212/32  *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                      AS path: 300 I
                    > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix
172.16.20.20/32     *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                      AS path: I
                    > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix
172.16.100.0/24       *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
```


Copyright © 2017, Juniper Networks, Inc.

```
                      AS path: I
                    > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.204:10:10.255.245.212/32
                    *[BGP/170] 00:06:40, localpref 100, from 10.255.245.204
                      AS path: 300 I
                     > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:172.16.20.20/32
                    *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
                      AS path: I
                     > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:100.1.4.0/24
                    *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
                      AS path: I
                     > to 172.16.100.1 via ge-1/1/0.0, label-switched-path to_fix

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

instance1.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

show route damping

List of Syntax
Syntax on page 218
Syntax (EX Series Switch and QFX Series) on page 218

Syntax
show route damping (decayed | history | suppressed)
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch and QFX Series)
show route damping (decayed | history | suppressed)
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Display the BGP routes for which updates might have been reduced because of route flap damping.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

decayed—Display route damping entries that might no longer be valid, but are not suppressed.

history—Display entries that have already been withdrawn, but have been logged.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

suppressed—Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols.

Required Privilege Level
view

Related Documentation
- *clear bgp damping*
- *show policy damping*

List of Sample Output
show route damping decayed detail on page 221
show route damping history on page 222
show route damping history detail on page 222

Output Fields
Table 18 on page 219 lists the output fields for the **show route damping** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 18: show route damping Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table—for example, **inet.0**.	All levels
destinations	Number of destinations for which there are routes in the routing table.	All levels
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** • **holddown** (routes that are in a pending state before being declared inactive) • **hidden** (the routes are not used because of a routing policy)	All levels
destination-prefix **(entry, announced)**	Destination prefix. The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination.	**detail extensive**
[*protocol***,** *preference***]**	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.	All levels
Next-hop reference count	Number of references made to the next hop.	**detail extensive**
Source	IP address of the route source.	**detail extensive**
Next hop	Network layer address of the directly reachable neighboring system.	**detail extensive**
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word **Selected**.	**detail extensive**
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.	**detail extensive**
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.	**detail extensive**
State	Flags for this route. For a description of possible values for this field, see the output field table for the show route detail command.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 18: show route damping Output Fields *(continued)*

Field Name	Field Description	Level of Output
Local AS	AS number of the local routing device.	**detail extensive**
Peer AS	AS number of the peer routing device.	**detail extensive**
Age	How long the route has been known.	**detail extensive**
Metric	Metric for the route.	**detail extensive**
Task	Name of the protocol that has added the route.	**detail extensive**
Announcement bits	List of protocols that announce this route. ***n*-Resolve inet** indicates that the route is used for route resolution for next hops found in the routing table. ***n*** is an index used by Juniper Networks customer support only.	**detail extensive**
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.	All levels
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	**brief** none
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word **Selected**.	**brief** none
Communities	Community path attribute for the route. See the output field table for the show route detail command.	**detail extensive**
Localpref	Local preference value included in the route.	All levels
Router ID	BGP router ID as advertised by the neighbor in the open message.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 18: show route damping Output Fields *(continued)*

Field Name	Field Description	Level of Output
Merit (last update/now)	Last updated and current figure-of-merit value.	**detail extensive**
damping-parameters	Name that identifies the damping parameters used, which is defined in the damping statement at the **[edit policy-options]** hierarchy level.	**detail extensive**
Last update	Time of most recent change in path attributes.	**detail extensive**
First update	Time of first change in path attributes, which started the route damping process.	**detail extensive**
Flaps	Number of times the route has gone up or down or its path attributes have changed.	**detail extensive**
Suppressed	(**suppressed** keyword only) This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it.	All levels
Reusable in	(**suppressed** keyword only) Time when a suppressed route will again be available.	All levels
Preference will be	(**suppressed** keyword only) Preference value that will be applied to the route when it is again active.	All levels

Sample Output

show route damping decayed detail

```
user@host> show route damping decayed detail
inet.0: 173319 destinations, 1533668 routes (172625 active, 4 holddown, 108083
hidden)
10.0.111.0/24 (7 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 151973
                Source: 172.23.2.129
                Next hop: via so-1/2/0.0
                Next hop: via so-5/1/0.0, selected
                Next hop: via so-6/0/0.0
                Protocol next hop: 172.23.2.129
                Indirect next hop: 89a1a00 264185
                State: <Active Ext>
                Local AS: 64500 Peer AS:   64490
                Age: 3:28       Metric2: 0
                Task: BGP_64490.172.23.2.129+179
                Announcement bits (6): 0-KRT 1-RT 4-KRT 5-BGP.0.0.0.0+179

               6-Resolve tree 2 7-Resolve tree 3
                AS path: 64499 64510 645511 645511 645511 645511 I ()
                Communities: 65551:390 65551:2000 65551:3000 65550:701
                Localpref: 100
                Router ID: 172.23.2.129
                Merit (last update/now): 1934/1790
                damping-parameters: damping-high
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Last update:       00:03:28 First update:       00:06:40
                    Flaps: 2
```

show route damping history

```
user@host> show route damping history
inet.0: 173320 destinations, 1533529 routes (172624 active, 6 holddown, 108122
hidden)
+ = Active Route, - = Last Active, * = Both

10.108.0.0/15       [BGP ] 2d 22:47:58, localpref 100
                      AS path: 64220 65541 65542 I
                    > to 192.168.60.85 via so-3/1/0.0
```

show route damping history detail

```
user@host> show route damping history detail
inet.0: 173319 destinations, 1533435 routes (172627 active, 2 holddown, 108105
hidden)
10.108.0.0/15 (3 entries, 1 announced)
         BGP                 /-101
                Next-hop reference count: 69058
                Source: 192.168.60.85
                Next hop: 192.168.60.85 via so-3/1/0.0, selected
                State: <Hidden Ext>
                Inactive reason: Unusable path
                Local AS: 64500 Peer AS:  64220
                Age: 2d 22:48:10
                Task: BGP_64220.192.168.60.85+179
                AS path: 64220 65541 65542 I ()
                Communities: 65541:390 65541:2000 65541:3000 65504:3561
                Localpref: 100
                Router ID: 192.168.80.25
                Merit (last update/now): 1000/932
                damping-parameters: set-normal
                Last update:       00:01:05 First update:       00:01:05
                Flaps: 1
```


Copyright © 2017, Juniper Networks, Inc.

show route detail

List of Syntax
Syntax on page 223
Syntax (EX Series Switches) on page 223

Syntax
show route detail
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route detail
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Comand introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Display detailed information about the active entries in the routing tables.

Options
none—Display all active entries in the routing table on all systems.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route detail on page 234
show route detail (with BGP Multipath) on page 240
show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 241
show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 241

Output Fields
Table 19 on page 223 describes the output fields for the **show route detail** command. Output fields are listed in the approximate order in which they appear.

Table 19: show route detail Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)

Copyright © 2017, Juniper Networks, Inc.

Table 19: show route detail Output Fields *(continued)*

Field Name	Field Description

route-destination **(entry, announced)**	Route destination (for example:10.0.0.1/24). The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• ***MPLS-label*** (for example, 80001).
	• ***interface-name*** (for example, ge-1/0/2).
	• ***neighbor-address*****:*****control-word-status*****:*****encapsulation type*****:*****vc-id*****:*****source*** (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.

label stacking	**(**Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).

Copyright © 2017, Juniper Networks, Inc.

Table 19: show route detail Output Fields *(continued)*

Field Name	Field Description
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • ***—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value. **Preference2** values are signed integers, that is, **Preference2** values can be either positive or negative values. However, Junos OS evaluates **Preference2** values as unsigned integers that are represented by positive values. Based on the **Preference2** values, Junos OS evaluates a preferred route differently in the following scenarios: • **Both Signed Preference2 values** ◦ Route A = -101 ◦ Route B = -156 Where both the **Preference2** values are signed, Junos OS evaluates only the unsigned value of **Preference2** and Route A, which has a lower **Preference2** value is preferred. • **Unsigned Preference2 values** Now consider both unsigned **Preference2** values: ◦ Route A = 4294967096 ◦ Route B = 200 Here, Junos OS considers the lesser **Preference2** value and Route B with a **Preference2** value of 200 is preferred because it is less than 4294967096. • **Combination of signed and unsigned Preference2 values** When **Preference2** values of two routes are compared, and for one route the **Preference2** is a signed value, and for the other route it is an unsigned value, Junos OS prefers the route with the positive **Preference2** value over the negative **Preference2** value. For example, consider the following signed and unsigned **Preference2** values: ◦ Route A = -200 ◦ Route B = 200 In this case, Route B with a **Preference2** value of 200 is preferred although this value is greater than -200, because Junos OS evaluates only the unsigned value of the **Preference2** value.
Level	(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).

Copyright © 2017, Juniper Networks, Inc.

Table 19: show route detail Output Fields *(continued)*

Field Name	Field Description
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 20 on page 229.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 21 on page 231.
Local AS	AS number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.

Copyright © 2017, Juniper Networks, Inc.

Table 19: show route detail Output Fields *(continued)*

Field Name	Field Description
Metric*n*	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see show route table.
Task	Name of the protocol that has added the route.
Announcement bits	The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the KRT for installing the route into the Packet Forwarding Engine, to a resolve tree, a L2 VC, or even a VPN. For example, ***n*-Resolve inet** indicates that the specified route is used for route resolution for next hops found in the routing table. • ***n***—An index used by Juniper Networks customer support only.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **Recorded**—The AS path is recorded by the sample process (sampled). • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.

Copyright © 2017, Juniper Networks, Inc.

Table 19: show route detail Output Fields *(continued)*

Field Name	Field Description
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
ORR Generation-ID	Displays the optimal route reflection (ORR) generation identifier. ISIS and OSPF interior gateway protocol (IGP) updates filed whenever any of the corresponding ORR route has its metric valued changed, or if the ORR route is added or deleted.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 22 on page 233 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or **Site Down**.
mtu	Maximum transmission unit (MTU) information.

Copyright © 2017, Juniper Networks, Inc.

Table 19: show route detail Output Fields *(continued)*

Field Name	Field Description
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 20 on page 229 describes all possible values for the Next-hop Types output field.

Table 20: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.

Copyright © 2017, Juniper Networks, Inc.

Table 20: Next-hop Types Output Field Values *(continued)*

Next-Hop Type	Description
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrt)	Regular multicast next hop.
Router	A specific node or set of nodes to which the routing device forwards packets that match the route prefix. To qualify as next-hop type router, the route must meet the following criteria: • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.

Copyright © 2017, Juniper Networks, Inc.

Table 20: Next-hop Types Output Field Values *(continued)*

Next-Hop Type	Description
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 21 on page 231 describes all possible values for the State output field. A route can be in more than one state (for example, **<Active NoReadvrt Int Ext>**).

Table 21: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.

Copyright © 2017, Juniper Networks, Inc.

Table 21: State Output Field Values *(continued)*

Value	Description
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGP path is available.
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.
Programmed	Route installed programatically by on-box or off-box applications using API.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available
Router ID	Path through a neighbor with lower ID is available.

Copyright © 2017, Juniper Networks, Inc.

Table 21: State Output Field Values *(continued)*

Value	Description
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Table 22 on page 233 describes the possible values for the Communities output field.

Table 22: Communities Output Field Values

Value	Description
area-number	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is **0**. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: ***local AS number:link-bandwidth-number***	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
link-bandwidth-number	Link-bandwidth number: from **0** through **4,294,967,295** (bytes per second).
local AS number	Local AS number: from **1** through **65,535**.
options	1 byte. Currently this is only used if the route type is **5** or **7**. Setting the least significant bit in the field indicates that the route carries a type 2 metric.
origin	(Used with VPNs) Identifies where the route came from.
ospf-route-type	1 byte, encoded as **1** or **2** for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); **3** for summary routes; **5** for external routes (area number must be**0**); **7** for NSSA routes; or **129** for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x8000**. The format is ***area-number:ospf-route-type:options***.
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x0306**. The format is ***area-number:ospf-route-type:options***.

Copyright © 2017, Juniper Networks, Inc.

Table 22: Communities Output Field Values *(continued)*

Value	Description
target	Defines which VPN the route participates in; **target** has the format ***32-bit IP address:16-bit number***. For example, 10.19.0.0:100.
unknown IANA	Incoming IANA codes with a value between **0x1** and **0x7fff**. This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above **0x8000**. This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route detail

```
user@host> show route detail

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 1:31:43
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

10.31.1.0/30 (2 entries, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 2
                Next hop: via so-0/3/0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:30:17
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:    69
                Age: 1:30:17    Metric: 1
                ORR Generation-ID: 1
     Area: 0.0.0.0
                Task: OSPF
                AS path: I

10.31.1.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 7
                Interface: so-0/3/0.0
```


Copyright © 2017, Juniper Networks, Inc.

```
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:30:20
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:29:56    Metric: 2
                Area: 0.0.0.0
                ORR Generation-ID: 1
     Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

172.16.233.2/32 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

172.16.233.22/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:43
                Task: IGMP
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
                State: <Active Int>
                Local AS:    69
                Age: 1:25:49    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
```

Copyright © 2017, Juniper Networks, Inc.

```
                AS path: I

10.255.71.238/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r2
                State: <Active Int>
                Local AS:    69
                Age: 1:25:49    Metric: 1
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

private__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0   (1 entry, 1 announced)
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
                State: <Active Int>
                Local AS:    69
                Age: 1:31:45    Metric: 1
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

  299840 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 299840 /52 -> {indirect(1048575)}
          *RSVP   Preference: 7/2
                  Next hop type: Flood
                  Address: 0x9174a30
                  Next-hop reference count: 4
                  Next hop type: Router, Next hop index: 798
                  Address: 0x9174c28
                  Next-hop reference count: 2
                  Next hop: 172.16.0.2 via lt-1/2/0.9 weight 0x1
                  Label-switched-path R2-to-R4-2p2mp
                  Label operation: Pop
                  Next hop type: Router, Next hop index: 1048574
                  Address: 0x92544f0
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Next-hop reference count: 2
                    Next hop: 172.16.0.2 via lt-1/2/0.7 weight 0x1
                    Label-switched-path R2-to-R200-p2mp
                    Label operation: Pop
                    Next hop: 172.16.0.2 via lt-1/2/0.5 weight 0x8001
                    Label operation: Pop
                    State: <Active Int>
                    Age: 1:29       Metric: 1
                    Task: RSVP
                    Announcement bits (1): 0-KRT
                    AS path: I...

800010 (1 entry, 1 announced)
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: via vt-3/2/0.32769, selected
                Label operation: Pop
                State: <Active Int>
                Age: 1:29:30
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 800012, Push 100096(top)
                Protocol next hop: 10.255.70.103
                Push 800012
                Indirect next hop: 87272e4 1048574
                State: <Active Int>
                Age: 1:29:30    Metric2: 2
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:44
```

Copyright © 2017, Juniper Networks, Inc.

```
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:43
                Task: MLD
                Announcement bits (1): 0-KRT
                AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.16385, selected
                State: <Active NoReadvrt Int>
                Age: 1:31:44
                Task: IF
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 1:25:49    Metric2: 1
                AIGP 210
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Communities: target:11111:1 Layer2-info: encaps:VPLS,
                    control flags:, mtu: 0
                    Label-base: 800008, range: 8
                    Localpref: 100
                    Router ID: 10.255.70.103
                    Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
        *L2VPN  Preference: 170/-1
                    Next-hop reference count: 5
                    Protocol next hop: 10.255.71.52
                    Indirect next hop: 0 -
                    State: <Active Int Ext>
                    Age: 1:31:40    Metric2: 1
                    Task: green-l2vpn
                    Announcement bits (1): 1-BGP.0.0.0.0+179
                    AS path: I
                    Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
                    mtu: 0
                    Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
        *L2VPN  Preference: 170/-101
                    Next-hop reference count: 5
                    Protocol next hop: 10.255.71.52
                    Indirect next hop: 0 -
                    State: <Active Int Ext>
                    Age: 1:31:40    Metric2: 1
                    Task: green-l2vpn
                    Announcement bits (1): 1-BGP.0.0.0.0+179
                    AS path: I
                    Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
                    Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                    Next hop: via so-1/1/2.0 weight 1, selected
                    Label-switched-path my-lsp
                    Label operation: Push 100000[0]
                    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
                    State: <Active Int>
                    Local AS:    99
                    Age: 10:21
                    Task: l2 circuit
                    Announcement bits (1): 0-LDP
                    AS path: I
                    VC Label 100000, MTU 1500, VLAN ID 512

inet.0: 45 destinations, 47 routes (44 active, 0 holddown, 1 hidden)
1.1.1.3/32 (1 entry, 1 announced)
        *IS-IS  Preference: 18
                    Level: 2
                    Next hop type: Router, Next hop index: 580
                    Address: 0x9db6ed0
                    Next-hop reference count: 8
                    Next hop: 10.1.1.6 via lt-1/0/10.5, selected
                    Session Id: 0x18a
                    State: <Active Int>
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Local AS:     2
                    Age: 1:32       Metric: 10
                    Validation State: unverified
                    ORR Generation-ID: 1
                    Task: IS-IS
                    Announcement bits (3): 0-KRT 5-Resolve tree 4 6-Resolve_IGP_FRR
task
                    AS path: I

inet.0: 61 destinations, 77 routes (61 active, 1 holddown, 0 hidden)
1.1.1.1/32 (2 entries, 1 announced)
        *OSPF   Preference: 10
                Next hop type: Router, Next hop index: 673
                Address: 0xc008830
                Next-hop reference count: 3
                Next hop: 10.1.1.1 via ge-0/0/2.0, selected
                Session Id: 0x1b7
                State: <Active Int>
                Local AS:     1
                Age: 3:06:59    Metric: 100
                Validation State: unverified
                ORR Generation-ID: 1
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 1-KRT 9-Resolve tree 2
                AS path: I
```

show route detail (with BGP Multipath)

```
user@host> show route detail

10.1.1.8/30 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 262142
                Address: 0x901a010
                Next-hop reference count: 2
                Source: 10.1.1.2
                Next hop: 10.1.1.2 via ge-0/3/0.1, selected
                Next hop: 10.1.1.6 via ge-0/3/0.5
                State: <Active Ext>
                Local AS:     1 Peer AS:     2
                Age: 5:04:43
               Validation State: unverified
                Task: BGP_2.10.1.1.2+59955
                Announcement bits (1): 0-KRT
                AS path: 2 I
                Accepted Multipath
                Localpref: 100
                Router ID: 172.16.1.2
         BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 678
                Address: 0x8f97520
                Next-hop reference count: 9
                Source: 10.1.1.6
                Next hop: 10.1.1.6 via ge-0/3/0.5, selected
                State: <NotBest Ext>
                Inactive reason: Not Best in its group - Active preferred
                Local AS:     1 Peer AS:     2
                Age: 5:04:43
            Validation State: unverified
                Task: BGP_2.10.1.1.6+58198
```


Copyright © 2017, Juniper Networks, Inc.

```
                    AS path: 2 I
                    Accepted MultipathContrib
                    Localpref: 100
                    Router ID: 172.16.1.3
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
                Next hop: via vt-0/1/0.1
                Next-hop index: 661
                Label operation: Pop
                Address: 0x9172130
                Next hop: via so-0/0/3.0
                Next-hop index: 654
                Label operation: Swap 299872
                State: **Active Int>
                Local AS:  1001
                Age: 8:20       Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
               FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
 src 192.168.142.2
```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Address: 0x2735208
                Next-hop reference count: 3
                Next hop type: Router, Next hop index: 1397
                Address: 0x2735d2c
                Next-hop reference count: 3
                Next hop: 1.3.8.2 via ge-1/2/22.0
                Label operation: Pop
                Load balance label:  None;
                Next hop type: Router, Next hop index: 1395
                Address: 0x2736290
                Next-hop reference count: 3
                Next hop: 1.3.4.2 via ge-1/2/18.0
                Label operation: Pop
                Load balance label:  None;
                State: <Active Int AckRequest MulticastRPF>
                Local AS:    10
                Age: 54:05      Metric: 1
                Validation State: unverified
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                FECs bound to route: P2MP root-addr 172.16.1.1, grp: 232.1.1.1,
src: 192.168.219.11
```

Copyright © 2017, Juniper Networks, Inc.

```
Primary Upstream : 172.16.1.3:0--172.16.1.2:0
   RPF Nexthops  :
        ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
        ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
Backup Upstream : 172.16.1.3:0--172.16.1.6:0
   RPF Nexthops :
        ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xfffe
        ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xfffe
```


Copyright © 2017, Juniper Networks, Inc.

show route exact

List of Syntax
Syntax on page 243
Syntax (EX Series Switches) on page 243

Syntax
show route exact *destination-prefix*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route exact *destination-prefix*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display only the routes that exactly match the specified address or range of addresses.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

destination-prefix—Address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route exact on page 243
show route exact detail on page 243
show route exact extensive on page 244
show route exact terse on page 244

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route exact

```
user@host> show route exact 207.17.136.0/24

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24    *[Static/5] 2d 03:30:22
                    > to 192.168.71.254 via fxp0.0
```

show route exact detail

```
user@host> show route exact 207.17.136.0/24 detail

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

```
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2d 3:30:26
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I
```

show route exact extensive

```
user@host> show route exact 207.17.136.0/24 extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 1:25:18
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I
```

show route exact terse

```
user@host> show route exact 207.17.136.0/24 terse

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 207.17.136.0/24    S   5                      >192.168.71.254
```


Copyright © 2017, Juniper Networks, Inc.

show route export

List of Syntax
Syntax on page 245
Syntax (EX Series Switches) on page 245

Syntax
show route export
<brief | detail>
<instance <*instance-name*> | *routing-table-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route export
<brief | detail>
<instance <*instance-name*> | *routing-table-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances.

Options
none—(Same as **brief**.) Display standard information about policy-based export for all instances and routing tables on all systems.

brief | detail—(Optional) Display the specified level of output.

instance <*instance-name*>—(Optional) Display a particular routing instance for which policy-based export is currently enabled.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

routing-table-name—(Optional) Display information about policy-based export for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route export inet** command).

Required Privilege Level
view

List of Sample Output
show route export on page 246
show route export detail on page 246
show route export instance detail on page 246

Output Fields
Table 23 on page 245 lists the output fields for the **show route export** command. Output fields are listed in the approximate order in which they appear.

Table 23: show route export Output Fields

Field Name	Field Description	Level of Output
Table or *table-name*	Name of the routing tables that either import or export routes.	All levels
Routes	Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one.	**brief** none

Copyright © 2017, Juniper Networks, Inc.

Table 23: show route export Output Fields *(continued)*

Field Name	Field Description	Level of Output
Export	Whether the table is currently exporting routes to other tables: **Y** or **N** (Yes or No).	**brief** none
Import	Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.)	**detail**
Flags	(**instance** keyword only) Flags for this feature on this instance: • **config auto-policy**—The policy was deduced from the configured IGP export policies. • **cleanup**—Configuration information for this instance is no longer valid. • **config**—The instance was explicitly configured.	**detail**
Options	(**instance** keyword only) Configured option displays the type of routing tables the feature handles: • **unicast**—Indicates *instance*.inet.0. • **multicast**—Indicates *instance*.inet.2. • **unicast multicast**—Indicates *instance*.inet.0 and *instance*.inet.2.	**detail**
Import policy	(**instance** keyword only) Policy that **route export** uses to construct the import-export matrix. Not displayed if the instance type is **vrf**.	**detail**
Instance	(**instance** keyword only) Name of the routing instance.	**detail**
Type	(**instance** keyword only) Type of routing instance: **forwarding**, **non-forwarding**, or **vrf**.	**detail**

Sample Output

show route export

```
user@host> show route export
Table                            Export           Routes
inet.0                           N                     0
black.inet.0                     Y                     3
red.inet.0                       Y                     4
```

show route export detail

```
user@host>  show route export detail
inet.0                                          Routes:        0
black.inet.0                                    Routes:        3
  Import: [ inet.0 ]
red.inet.0                                      Routes:        4
  Import: [ inet.0 ]
```

show route export instance detail

```
user@host>  show route export instance detail
Instance: master                          Type: forwarding
  Flags: <config auto-policy> Options: <unicast multicast>
  Import policy: [ (ospf-master-from-red || isis-master-from-black) ]
```


Copyright © 2017, Juniper Networks, Inc.

```
Instance: black                          Type: non-forwarding
Instance: red                            Type: non-forwarding
```

Copyright © 2017, Juniper Networks, Inc.

show route extensive

List of Syntax
Syntax on page 248
Syntax (EX Series Switches) on page 248

Syntax
show route extensive
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route extensive
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display extensive information about the active entries in the routing tables.

Options
none—Display all active entries in the routing table.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route extensive on page 255
show route extensive (Access Route) on page 262
show route extensive (BGP PIC Edge) on page 262
show route extensive (FRR and LFA) on page 263
show route extensive (IS-IS) on page 264
show route extensive (Route Reflector) on page 264
show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 264
show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 265

Output Fields
Table 24 on page 248 describes the output fields for the **show route extensive** command. Output fields are listed in the approximate order in which they appear.

Table 24: show route extensive Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.

Copyright © 2017, Juniper Networks, Inc.

Table 24: show route extensive Output Fields *(continued)*

Field Name	Field Description
***number* routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active). • **holddown** (routes that are in the pending state before being declared inactive). • **hidden** (routes that are not used because of a routing policy).

***route-destination* (entry, announced)**	Route destination (for example: 10.0.0.1/24). The **entry** value is the number of route for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• ***MPLS-label*** (for example, 80001).
	• ***interface-name*** (for example, ge-1/0/2).
	• ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.

TSI	Protocol header information.
label stacking	(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.

Copyright © 2017, Juniper Networks, Inc.

Table 24: show route extensive Output Fields *(continued)*

Field Name	Field Description
Level	(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see the Output Field table in the show route detail command.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Offset	Whether the metric has been increased or decreased by an offset value.
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.

Copyright © 2017, Juniper Networks, Inc.

Table 24: show route extensive Output Fields *(continued)*

Field Name	Field Description
label-operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Indirect next hops	When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved. When BGP PIC Edge is enabled, the output lines that contain **Indirect next hop: weight** follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values: • 0x1 indicates active next hops. • 0x4000 indicates passive next hops.
State	State of the route (a route can be in more than one state). See the Output Field table in the show route detail command.
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).
Weight	Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown. For sample output, see show route table.

Copyright © 2017, Juniper Networks, Inc.

Table 24: show route extensive Output Fields *(continued)*

Field Name	Field Description
Inactive reason	If the route is inactive, the reason for its current state is indicated. Typical reasons include: • **Active preferred**—Currently active route was selected over this route. • **Always compare MED**—Path with a lower multiple exit discriminator (MED) is available. • **AS path**—Shorter AS path is available. • **Cisco Non-deterministic MED selection**—Cisco nondeterministic MED is enabled and a path with a lower MED is available. • **Cluster list length**—Path with a shorter cluster list length is available. • **Forwarding use only**—Path is only available for forwarding purposes. • **IGP metric**—Path through the next hop with a lower IGP metric is available. • **IGP metric type**—Path with a lower OSPF link-state advertisement type is available. • **Interior > Exterior > Exterior via Interior**—Direct, static, IGP, or EBGP path is available. • **Local preference**—Path with a higher local preference value is available. • **Next hop address**—Path with a lower metric next hop is available. • **No difference**—Path from a neighbor with a lower IP address is available. • **Not Best in its group**—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed). • **Number of gateways**—Path with a higher number of next hops is available. • **Origin**—Path with a lower origin code is available. • **OSPF version**—Path does not support the indicated OSPF version. • **RIB preference**—Route from a higher-numbered routing table is available. • **Route destinguisher**—64-bit prefix added to IP subnets to make them unique. • **Route metric or MED comparison**—Route with a lower metric or MED is available. • **Route preference**—Route with a lower preference value is available. • **Router ID**—Path through a neighbor with a lower ID is available. • **Unusable path**—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved. • **Update source**—Last tiebreaker is the lowest IP address value.
Local AS	Autonomous system (AS) number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see show route table.

Copyright © 2017, Juniper Networks, Inc.

Table 24: show route extensive Output Fields *(continued)*

Field Name	Field Description
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that are consumers of the route. Using the following output as an example, **Announcement bits (3): 0-KRT 5-Resolve tree 2 8-BGP RT Background** there are (3) announcement bits to reflect the three clients (protocols) that have state for this route: Kernel (0-KRT), 5 (resolution tree process 2), and 8 (BGP). The notation ***n*-Resolve inet** indicates that the route is used for route resolution for next hops found in the routing table. ***n*** is an index used by Juniper Networks customer support only.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **Recorded**—The AS path is recorded by the sample process (sampled). • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
AS path: I <Originator>	(For route reflected output only) Originator ID attribute set by the route reflector.

Copyright © 2017, Juniper Networks, Inc.

Table 24: show route extensive Output Fields *(continued)*

Field Name	Field Description
route status	Indicates the status of a BGP route: • **Accepted**—The specified BGP route is imported by the default BGP policy. • **Import**—The route is imported into a Layer 3 VPN routing instance. • **Import-Protect**—A remote instance egress that is protected. • **Multipath**—A BGP multipath active route. • **MultipathContrib**—The route is not active but contributes to the BGP multipath. • **Protect**—An egress route that is protected. • **Stale**—A route that is marked stale due to graceful restart.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.
Originator ID	(For route reflected output only) Address of router that originally sent the route to the route reflector.
Prefixes bound to route	Forwarding Equivalent Class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or Site Down.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.

Copyright © 2017, Juniper Networks, Inc.

Table 24: show route extensive Output Fields *(continued)*

Field Name	Field Description
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix.
Node path count	Number of nodes in the path.
Forwarding nexthops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route extensive

```
user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
203.0.113.10/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.10/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    64496
                Age: 1:34:06
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

203.0.113.30/30 (2 entries, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 2
                Next hop: via so-0/3/0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:32:40
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
```

Copyright © 2017, Juniper Networks, Inc.

```
                State: <Int>
                Inactive reason: Route Preference
                Local AS:    64496
                Age: 1:32:40    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

203.0.113.103/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 7
                Interface: so-0/3/0.0
                State: <Active NoReadvrt Int>
                Local AS:    644969
                Age: 1:32:43
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...

203.0.113.203/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.203/30 -> {203.0.113.216}
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 203.0.113.216 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:32:19    Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

198.51.100.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.2/32 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

198.51.100.22/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.22/32 -> {}
        *IGMP   Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:06
```

 Copyright © 2017, Juniper Networks, Inc.

```
                Task: IGMP
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

203.0.113.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
                State: <Active Int>
                Local AS:    64496
                Age: 1:28:12    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

203.0.113.238/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r2
                State: <Active Int>
                Local AS:    64496
                Age: 1:28:12    Metric: 1
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0   (1 entry, 1 announced)
TSI:
KRT in-kernel 0      /36 -> {}
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:08    Metric: 1
```

Copyright © 2017, Juniper Networks, Inc.

```
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
  299840 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 299840 /52 -> {indirect(1048575)}
          *RSVP   Preference: 7/2
                  Next hop type: Flood
                  Address: 0x9174a30
                  Next-hop reference count: 4
                  Next hop type: Router, Next hop index: 798
                  Address: 0x9174c28
                  Next-hop reference count: 2
                  Next hop: 198.51.100.2 via lt-1/2/0.9 weight 0x1
                  Label-switched-path R2-to-R4-2p2mp
                  Label operation: Pop
                  Next hop type: Router, Next hop index: 1048574
                  Address: 0x92544f0
                  Next-hop reference count: 2
                  Next hop: 198.51.100.2 via lt-1/2/0.7 weight 0x1
                  Label-switched-path R2-to-R200-p2mp
                  Label operation: Pop
                  Next hop: 198.51.100.2 via lt-1/2/0.5 weight 0x8001
                  Label operation: Pop
                  State: <Active Int>
                  Age: 1:29       Metric: 1
                  Task: RSVP
                  Announcement bits (1): 0-KRT
                  AS path: I...

800010 (1 entry, 1 announced)

TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: via vt-3/2/0.32769, selected
                Label operation: Pop
                State: <Active Int>
                Age: 1:31:53
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0      /16 -> {indirect(1048574)}
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 800012, Push 100096(top)
                Protocol next hop: 203.0.113.103
                Push 800012
                Indirect next hop: 87272e4 1048574
                State: <Active Int>
                Age: 1:31:53    Metric2: 2
```

 Copyright © 2017, Juniper Networks, Inc.

```
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Indirect next hops: 1
                        Protocol next hop: 203.0.113.103 Metric: 2
                        Push 800012
                        Indirect next hop: 87272e4 1048574
                        Indirect path forwarding next hops: 1
                               Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1

                        203.0.113.103/32 Originating RIB: inet.3
                          Metric: 2                     Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 203.0.113.216 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

2001:db8::10:255:71:52/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
```

Copyright © 2017, Juniper Networks, Inc.

```
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:06
                Task: MLD
                Announcement bits (1): 0-KRT
                AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.16385, selected
                State: <Active NoReadvrt Int>
                Age: 1:34:07
                Task: IF
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

203.0.113.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 203.0.113.103:1
                Next-hop reference count: 7
                Source: 203.0.113.103
                Protocol next hop: 203.0.113.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    64496 Peer AS:    64496
                Age: 1:28:12    Metric2: 1
                Task: BGP_69.203.0.113.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 203.0.113.103
                Primary Routing Table bgp.l2vpn.0

203.0.113.152:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
        *L2VPN  Preference: 170/-1
                Next-hop reference count: 5
                Protocol next hop: 203.0.113.152
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
```

 Copyright © 2017, Juniper Networks, Inc.

```
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8, status-vector: 0x9F

203.0.113.152:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
        *L2VPN  Preference: 170/-101
                Next-hop reference count: 5
                Protocol next hop: 203.0.113.152
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
                Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

203.0.113.163:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 203.0.113.163 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    64499
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512


203.0.113.55/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
  198.51.100.0/24 -> {Push 300112}
        *BGP    Preference: 170/-101
                Next hop type: Router
                Address: 0x925c208
                Next-hop reference count: 2
                Source: 203.0.113.9
                Next hop: 203.0.113.9 via ge-1/2/0.15, selected
                Label operation: Push 300112
                Label TTL action: prop-ttl
                State: <Active Ext>
                Local AS:  64509 Peer AS: 65539
                Age: 1w0d 23:06:56
                AIGP: 25
                Task: BGP_65539.203.0.113.9+56732
                Announcement bits (1): 0-KRT
                AS path: 65539 64508 I
                Accepted
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Route Label: 300112
                    Localpref: 100
                    Router ID: 213.0.113.99
```

show route extensive (Access Route)

```
user@host> show route 203.0.113.102 extensive
inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
203.0.113.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.102/32 -> {192.0.2.2}
OSPF area : 0.0.0.0, LSA ID : 203.0.113.102, LSA type : Extern
        *Access Preference: 13
                Next-hop reference count: 78472
                Next hop: 192.0.2.2 via fe-0/0/0.0, selected
                State: <Active Int>
       Age: 12
                Task: RPD Unix Domain Server./var/run/rpd_serv.local
                Announcement bits (2): 0-KRT 1-OSPFv2
                AS path: I

user@host> show route 2001:db8:4641:1::/48 extensive

inet6.0: 75 destinations, 81 routes (75 active, 0 holddown, 0 hidden)
  2001:db8:4641:1::/48 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 2001:db8:4641:1::/48 -> {#0 0.13.1.0.0.1}
          *Access Preference: 13
                  Next hop type: Router, Next hop index: 74548
                  Address: 0x1638c1d8
                  Next-hop reference count: 6
                  Next hop: #0 0.13.1.0.0.1 via demux0.1073753267, selected
                  Session Id: 0x0
                  State: <Active Int>
                  Age: 4:17
                  Validation State: unverified
                  Task: RPD Unix Domain Server./var/run/rpd_serv.local
                  Announcement bits (2): 0-KRT 4-Resolve tree 2
                  AS path: I
2001:db8:4641:1::/128 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 2001:db8:4641:1::/128 -> {#0 0.13.1.0.0.1}
          *Access-internal Preference: 12
                  Next hop type: Router, Next hop index: 74548
                  Address: 0x1638c1d8
                  Next-hop reference count: 6
                  Next hop: #0 0.13.1.0.0.1 via demux0.1073753267, selected
                  Session Id: 0x0
                  State: <Active Int>
                  Age: 4:17
                  Validation State: unverified
                  Task: RPD Unix Domain Server./var/run/rpd_serv.local
                  Announcement bits (2): 0-KRT 4-Resolve tree 2
                  AS path: I
```

show route extensive (BGP PIC Edge)

```
user@host> show route 198.51.100.6 extensive
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  198.51.100.6/32 (3 entries, 2 announced)
          State: <CalcForwarding>
```


Copyright © 2017, Juniper Networks, Inc.

```
  TSI:
  KRT in-kernel 198.51.100.6/32 -> {indirect(1048574), indirect(1048577)}
  Page 0 idx 0 Type 1 val 9219e30
      Nexthop: Self
      AS path: [2] 3 I
      Communities: target:2:1
  Path 198.51.100.6 from 198.51.100.4 Vector len 4.  Val: 0
..
           #Multipath Preference: 255
                  Next hop type: Indirect
                  Address: 0x93f4010
                  Next-hop reference count: 2
..
                  Protocol next hop: 198.51.1001.4
                  Push 299824
                  Indirect next hop: 944c000 1048574 INH Session ID: 0x3
                  Indirect next hop: weight 0x1
                  Protocol next hop: 198.51.100.5
                  Push 299824
                  Indirect next hop: 944c1d8 1048577 INH Session ID: 0x4
                  Indirect next hop: weight 0x4000
                  State: <ForwardingOnly Int Ext)>
                  Inactive reason: Forwarding use only
                  Age: 25         Metric2: 15
                  Validation State: unverified
                  Task: RT
                  Announcement bits (1): 0-KRT
                  AS path: 3 I
                  Communities: target:2:1
```

show route extensive (FRR and LFA)

```
user@host> show route 203.0.113.20 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
203.0.113.20/24 (2 entries, 1 announced)
        State: FlashAll
TSI:
KRT in-kernel 203.0.113.20/24 -> {Push 299776, Push 299792}
        *RSVP   Preference: 7/1
                Next hop type: Router, Next hop index: 1048574
                Address: 0xbbbc010
                Next-hop reference count: 5
                Next hop: 203.0.113.112 via ge-2/1/8.0 weight 0x1, selected
                Label-switched-path europa-d-to-europa-e
                Label operation: Push 299776
                Label TTL action: prop-ttl
                Session Id: 0x201
                Next hop: 203.0.113.122 via ge-2/1/4.0 weight 0x4001
                Label-switched-path europa-d-to-europa-e
                Label operation: Push 299792
                Label TTL action: prop-ttl
                Session Id: 0x202
                State: Active Int
                Local AS:   64500
                Age: 5:31  Metric: 2
                Task: RSVP
                Announcement bits (1): 0-KRT
                AS path: I
         OSPF   Preference: 10
                Next hop type: Router, Next hop index: 615
                Address: 0xb9d78c4
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Next-hop reference count: 7
                    Next hop: 203.0.113.112 via ge-2/1/8.0, selected
                    Session Id: 0x201
                    State: Int
                    Inactive reason: Route Preference
                    Local AS:   64500
                    Age: 5:35  Metric: 3
                    Area: 0.0.0.0
                    Task: OSPF
                    AS path: I
```

show route extensive (IS-IS)

```
user@host> show route extensive
IS-IS Preference: 15
                    Level: 1
                    Next hop type: Router, Next hop index: 1048577
                    Address: 0xXXXXXXXXXX
                    Next-hop reference count: YY
                    Next hop: 203.0.113.22 via ae1.0 balance 43%, selected
                    Session Id: 0x141
                    Next hop: 203.0.113.22 via ae0.0 balance 57%
```

show route extensive (Route Reflector)

```
user@host> show route extensive
203.0.113.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 203.0.113.0/8 -> {indirect(40)}
        *BGP    Preference: 170/-101
                Source: 192.168.4.214
                Protocol next hop: 198.51.100.192 Indirect next hop: 84ac908 40
                State: <Active Int Ext>
                Local AS: 65548 Peer AS: 65548
                Age: 3:09       Metric: 0       Metric2: 0
                Task: BGP_65548.192.168.4.214+1033
                Announcement bits (2): 0-KRT 4-Resolve inet.0
                AS path: 65544 64507 I <Originator>
                Cluster list:  198.51.100.1
                Originator ID: 203.0.113.88
                Communities: 7777:7777
                Localpref: 100
                Router ID: 203.0.113.4
                Indirect next hops: 1
                        Protocol next hop: 203.0.113.192 Metric: 0
                        Indirect next hop: 84ac908 40
                        Indirect path forwarding next hops: 0
                                Next hop type: Discard
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
                Next hop: via vt-0/1/0.1
                Next-hop index: 661
                Label operation: Pop
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Address: 0x9172130
                    Next hop: via so-0/0/3.0
                    Next-hop index: 654
                    Label operation: Swap 299872
                    State: **Active Int>
                    Local AS:  64511
                    Age: 8:20       Metric: 1
                    Task: LDP
                    Announcement bits (1): 0-KRT
                    AS path: I
                   FECs bound to route: P2MP root-addr 203.0.113.166, grp 203.0.1.1,
 src 192.168.142.2
```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Address: 0x2735208
                Next-hop reference count: 3
                Next hop type: Router, Next hop index: 1397
                Address: 0x2735d2c
                Next-hop reference count: 3
                Next hop: 203.0.113.82 via ge-1/2/22.0
                Label operation: Pop
                Load balance label:  None;
                Next hop type: Router, Next hop index: 1395
                Address: 0x2736290
                Next-hop reference count: 3
                Next hop: 203.0.113.2 via ge-1/2/18.0
                Label operation: Pop
                Load balance label:  None;
                State: <Active Int AckRequest MulticastRPF>
                Local AS:    64500
                Age: 54:05      Metric: 1
                Validation State: unverified
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
               FECs bound to route: P2MP root-addr 198.51.100.1, grp: 232.1.1.1,
 src: 192.168.219.11
                Primary Upstream : 198.51.100.3:0--198.51.100.2:0
                   RPF Nexthops  :
                        ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
                        ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
                Backup Upstream : 198.51.100.3:0--198.51.100.6:0
                   RPF Nexthops :
                       ge-1/2/20.0, 198.51.100.96, Label: 301584, weight: 0xfffe

                       ge-1/2/19.0, 198.51.100.36, Label: 301584, weight: 0xfffe
```

Copyright © 2017, Juniper Networks, Inc.

show route flow validation

List of Syntax
Syntax on page 266
Syntax (EX Series Switches) on page 266

Syntax
show route flow validation
<brief | detail>
<*ip-prefix*>
<table *table-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route flow validation
<brief | detail>
<*ip-prefix*>
<table *table-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display flow route information.

Options
none—Display flow route information.

brief | detail—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

ip-prefix—(Optional) IP address for the flow route.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table *table-name*—(Optional) Display flow route information for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route flow validation inet** command).

Required Privilege Level
view

List of Sample Output
show route flow validation on page 267
show route flow validation (IPv6) on page 267

Output Fields
Table 25 on page 266 lists the output fields for the **show route flow validation** command. Output fields are listed in the approximate order in which they appear.

Table 25: show route flow validation Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table (for example, inet.0).	All levels
prefix	Route address.	All levels

Copyright © 2017, Juniper Networks, Inc.

Table 25: show route flow validation Output Fields *(continued)*

Field Name	Field Description	Level of Output
Active unicast route	Active route in the routing table.	All levels
Dependent flow destinations	Number of flows for which there are routes in the routing table.	All levels
Origin	Source of the route flow.	All levels
Neighbor AS	Autonomous system identifier of the neighbor.	All levels
Flow destination	Number of entries and number of destinations that match the route flow.	All levels
Unicast best match	Destination that is the best match for the route flow.	All levels
Flags	Information about the route flow.	All levels

Sample Output

show route flow validation

```
user@host> show route flow validation
inet.0:
10.0.5.0/24Active unicast route
Dependent flow destinations: 1
Origin: 192.168.224.218, Neighbor AS: 64501
Flow destination (3 entries, 1 match origin)
Unicast best match: 10.0.5.0/24
Flags: SubtreeApex Consistent
```

show route flow validation (IPv6)

```
user@host> show route flow validation
inet6.0:
2001:db8::11:11:11:0/120
                Active unicast route
                    Dependent flow destinations: 2
                    Origin: 2001:db8::13:14:2:2, Neighbor AS: 2000
2001:db8::11:11:11:10/128
                Flow destination (1 entries, 1 match origin, next-as)
                    Unicast best match: 2001:db8::11:11:11:0/120
                    Flags: Consistent
2001:db8::11:11:11:30/128
                Flow destination (1 entries, 1 match origin, next-as)
                    Unicast best match: 2001:db8::11:11:11:0/120
                    Flags: Consistent
```

Copyright © 2017, Juniper Networks, Inc.

show route forwarding-table

List of Syntax

Syntax on page 268
Syntax (MX Series Routers) on page 268
Syntax (TX Matrix and TX Matrix Plus Routers) on page 268

Syntax

show route forwarding-table
<detail | extensive | summary>
<all>
<ccc *interface-name*>
<destination *destination-prefix*>
<family *family* | matching *matching*>
<interface-name *interface-name*>
<label *name*>
<matching *matching*>
<multicast>
<table (default | *logical-system-name/routing-instance-name* | *routing-instance-name*)>
<vlan (all | *vlan-name*)>
<vpn *vpn*>

Syntax (MX Series Routers)

show route forwarding-table
<detail | extensive | summary>
<all>
<bridge-domain (all | *domain-name*)>
<ccc *interface-name*>
<destination *destination-prefix*>
<family *family* | matching *matching*>
<interface-name *interface-name*>
<label *name*>
<learning-vlan-id *learning-vlan-id*>
<matching *matching*>
<multicast>
<table (default | *logical-system-name/routing-instance-name* | *routing-instance-name*)>
<vlan (all | *vlan-name*)>
<vpn *vpn*>

Syntax (TX Matrix and TX Matrix Plus Routers)

show route forwarding-table
<detail | extensive | summary>
<all>
<ccc *interface-name*>
<destination *destination-prefix*>
<family *family* | matching *matching*>
<interface-name *interface-name*>
<matching *matching*>
<label *name*>
<lcc *number*>
<multicast>
<table *routing-instance-name*>
<vpn *vpn*>

Release Information

Command introduced before Junos OS Release 7.4.
Option **bridge-domain** introduced in Junos OS Release 7.5
Option **learning-vlan-id** introduced in Junos OS Release 8.4

Copyright © 2017, Juniper Networks, Inc.

Options **all** and **vlan** introduced in Junos OS Release 9.6.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.

NOTE: The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

Options **none**—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

detail | extensive | summary—(Optional) Display the specified level of output.

all—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

bridge-domain (all | *bridge-domain-name*)—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

ccc *interface-name*—(Optional) Display route entries for the specified circuit cross-connect interface.

destination *destination-prefix*—(Optional) Destination prefix.

family *family*—(Optional) Display routing table entries for the specified family: **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

interface-name *interface-name*—(Optional) Display routing table entries for the specified interface.

label *name*—(Optional) Display route entries for the specified label.

lcc *number*—(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers, display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

Copyright © 2017, Juniper Networks, Inc.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

learning-vlan-id *learning-vlan-id*—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

matching *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast—(Optional) Display routing table entries for multicast routes.

table (default | *logical-system-name/routing-instance-name* | *routing-instance-name*)—(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the show route instance command.

vlan (all | *vlan-name*)—(Optional) Display information for all VLANs or for the specified VLAN.

vpn *vpn*—(Optional) Display routing table entries for a specified VPN.

Required Privilege Level view

List of Sample Output

show route forwarding-table on page 273
show route forwarding-table detail on page 274
show route forwarding-table destination extensive (Weights and Balances) on page 274
show route forwarding-table extensive on page 275
show route forwarding-table extensive (RPF) on page 276
show route forwarding-table family mpls on page 277
show route forwarding-table family vpls on page 277
show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled) on page 277
show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics) on page 278
show route forwarding-table family vpls extensive on page 278
show route forwarding-table table default on page 279
show route forwarding-table table logical-system-name/routing-instance-name on page 280

Copyright © 2017, Juniper Networks, Inc.

show route forwarding-table vpn on page 281

Output Fields Table 26 on page 271 lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

Table 26: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Logical system	Name of the logical system. This field is displayed if you specify the **table** ***logical-system-name/routing-instance-name*** option on a device that is configured for and supports logical systems.	All levels
Routing table	Name of the routing table (for example, inet, inet6, mpls).	All levels
Address family	Address family (for example, **IP**, **IPv6**, **ISO**, **MPLS**, and **VPLS**).	All levels
Destination	Destination of the route.	**detail extensive**
Route Type (Type)	How the route was placed into the forwarding table. When the **detail** keyword is used, the route type might be abbreviated (as shown in parentheses): • **cloned** (**clon**)—(TCP or multicast only) Cloned route. • **destination** (**dest**)—Remote addresses directly reachable through an interface. • **destination down** (**iddn**)—Destination route for which the interface is unreachable. • **interface cloned** (**ifcl**)—Cloned route for which the interface is unreachable. • **route down** (**ifdn**)—Interface route for which the interface is unreachable. • **ignore** (**ignr**)—Ignore this route. • **interface** (**intf**)—Installed as a result of configuring an interface. • **permanent** (**perm**)—Routes installed by the kernel when the routing table is initialized. • **user**—Routes installed by the routing protocol process or as a result of the configuration.	All levels
Route Reference (RtRef)	Number of routes to reference.	**detail extensive**
Flags	Route type flags: • **none**—No flags are enabled. • **accounting**—Route has accounting enabled. • **cached**—Cache route. • **incoming-iface** ***interface-number***—Check against incoming interface. • **prefix load balance**—Load balancing is enabled for this prefix. • **rt nh decoupled**—Route has been decoupled from the next hop to the destination. • **sent to PFE**—Route has been sent to the Packet Forwarding Engine. • **static**—Static route.	**extensive**
Next hop	IP address of the next hop to the destination.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 26: show route forwarding-table Output Fields *(continued)*

Field Name	Field Description	Level of Output
Next hop Type (Type)	Next-hop type. When the **detail** keyword is used, the next-hop type might be abbreviated (as indicated in parentheses): • **broadcast (bcst)**—Broadcast. • **deny**—Deny. • **discard (dscd)** —Discard. • **hold**—Next hop is waiting to be resolved into a unicast or multicast type. • **indexed (idxd)**—Indexed next hop. • **indirect (indr)**—Indirect next hop. • **local (locl)**—Local address on an interface. • **routed multicast (mcrt)**—Regular multicast next hop. • **multicast (mcst)**—Wire multicast next hop (limited to the LAN). • **multicast discard (mdsc)**—Multicast discard. • **multicast group (mgrp)**—Multicast group member. • **receive (recv)**—Receive. • **reject (rjct)**—Discard. An ICMP unreachable message was sent. • **resolve (rslv)**—Resolving the next hop. • **unicast (ucst)**—Unicast. • **unilist (ulst)**—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.	**detail extensive**
Index	Software index of the next hop that is used to route the traffic for a given prefix.	**detail extensive** none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	**extensive**
Reference (NhRef)	Number of routes that refer to this next hop.	**detail extensive** none
Next-hop interface (Netif)	Interface used to reach the next hop.	**detail extensive** none
Weight	Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the **Balance** field description).	**extensive**
Balance	Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.	**extensive**
RPF interface	List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when **rpf-check** is configured on the interface.	**extensive**

Copyright © 2017, Juniper Networks, Inc.

Sample Output

show route forwarding-table

```
user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    46     4
0.0.0.0/32         perm     0                    dscd    44     1
172.16.1.0/24         ifdn     0                    rslv   608     1 ge-2/0/1.0
172.16.1.0/32         iddn     0 172.16.1.0           recv   606     1 ge-2/0/1.0
172.16.1.1/32         user     0                    rjct    46     4
172.16.1.1/32         intf     0 172.16.1.1           locl   607     2
172.16.1.1/32         iddn     0 172.16.1.1           locl   607     2
172.16.1.255/32       iddn     0 ff:ff:ff:ff:ff:ff  bcst   605     1 ge-2/0/1.0
10.0.0.0/24        intf     0                    rslv   616     1 ge-2/0/0.0
10.0.0.0/32        dest     0 10.0.0.0           recv   614     1 ge-2/0/0.0
10.0.0.1/32        intf     0 10.0.0.1           locl   615     2
10.0.0.1/32        dest     0 10.0.0.1           locl   615     2
10.0.0.255/32      dest     0 10.0.0.255         bcst   613     1 ge-2/0/0.0
10.1.1.0/24        ifdn     0                    rslv   612     1 ge-2/0/1.0
10.1.1.0/32        iddn     0 10.1.1.0           recv   610     1 ge-2/0/1.0
10.1.1.1/32        user     0                    rjct    46     4
10.1.1.1/32        intf     0 10.1.1.1           locl   611     2
10.1.1.1/32        iddn     0 10.1.1.1           locl   611     2
10.1.1.255/32      iddn     0 ff:ff:ff:ff:ff:ff  bcst   609     1 ge-2/0/1.0
10.206.0.0/16      user     0 10.209.63.254      ucst   419    20 fxp0.0
10.209.0.0/16      user     1 0:12:1e:ca:98:0    ucst   419    20 fxp0.0
10.209.0.0/18      intf     0                    rslv   418     1 fxp0.0
10.209.0.0/32      dest     0 10.209.0.0         recv   416     1 fxp0.0
10.209.2.131/32    intf     0 10.209.2.131       locl   417     2
10.209.2.131/32    dest     0 10.209.2.131       locl   417     2
10.209.17.55/32    dest     0 0:30:48:5b:78:d2   ucst   435     1 fxp0.0
10.209.63.42/32    dest     0 0:23:7d:58:92:ca   ucst   434     1 fxp0.0
10.209.63.254/32   dest     0 0:12:1e:ca:98:0    ucst   419    20 fxp0.0
10.209.63.255/32   dest     0 10.209.63.255      bcst   415     1 fxp0.0
10.227.0.0/16      user     0 10.209.63.254      ucst   419    20 fxp0.0

...

Routing table: iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                   rjct    27     1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf     0                   locl    28     1

Routing table: inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                   rjct     6     1
ff00::/8           perm     0                   mdsc     4     1
ff02::1/128        perm     0 ff02::1           mcst     3     1

Routing table: ccc
MPLS:
Interface.Label    Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct 16     1
100004(top)fe-0/0/1.0
```

Copyright © 2017, Juniper Networks, Inc.

show route forwarding-table detail

```
user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            user     2 0:90:69:8e:b1:1b   ucst   132     4 fxp0.0
default            perm     0                    rjct    14     1
10.1.1.0/24        intf     0 ff.3.0.21          ucst   322     1 so-5/3/0.0
10.1.1.0/32        dest     0 10.1.1.0           recv   324     1 so-5/3/0.0
10.1.1.1/32        intf     0 10.1.1.1           locl   321     1
10.1.1.255/32      dest     0 10.1.1.255         bcst   323     1 so-5/3/0.0
10.21.21.0/24      intf     0 ff.3.0.21          ucst   326     1 so-5/3/0.0
10.21.21.0/32      dest     0 10.21.21.0         recv   328     1 so-5/3/0.0
10.21.21.1/32      intf     0 10.21.21.1         locl   325     1
10.21.21.255/32    dest     0 10.21.21.255       bcst   327     1 so-5/3/0.0
127.0.0.1/32       intf     0 127.0.0.1          locl   320     1
172.17.28.19/32    clon     1 192.168.4.254      ucst   132     4 fxp0.0
172.17.28.44/32    clon     1 192.168.4.254      ucst   132     4 fxp0.0

...

Routing table: private1__.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    46     1
10.0.0.0/8         intf     0                    rslv   136     1 fxp1.0
10.0.0.0/32        dest     0 10.0.0.0           recv   134     1 fxp1.0
10.0.0.4/32        intf     0 10.0.0.4           locl   135     2
10.0.0.4/32        dest     0 10.0.0.4           locl   135     2

...

Routing table: iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    38     1

Routing table: inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    22     1
ff00::/8           perm     0                    mdsc    21     1
ff02::1/128        perm     0 ff02::1            mcst    17     1

...

Routing table: mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct 28     1
```

show route forwarding-table destination extensive (Weights and Balances)

```
user@host> show route forwarding-table destination 3.4.2.1 extensive
Routing table: inet [Index 0]
Internet:

Destination:  3.4.2.1/32
  Route type: user
  Route reference: 0                    Route interface-index: 0
```


Copyright © 2017, Juniper Networks, Inc.

```
  Flags: sent to PFE
  Next-hop type: unilist              Index: 262143   Reference: 1
  Nexthop: 172.16.4.4
  Next-hop type: unicast              Index: 335      Reference: 2
  Next-hop interface: so-1/1/0.0      Weight: 22      Balance: 3
  Nexthop: 145.12.1.2
  Next-hop type: unicast              Index: 337      Reference: 2
  Next-hop interface: so-0/1/2.0      Weight: 33      Balance: 33
```

show route forwarding-table extensive

```
user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:

Destination:  default
  Route type: user
  Route reference: 2                  Route interface-index: 0
  Flags: sent to PFE
  Nexthop: 00:00:5E:00:53:1b
  Next-hop type: unicast              Index: 132      Reference: 4
  Next-hop interface: fxp0.0

Destination:  default
  Route type: permanent
  Route reference: 0                  Route interface-index: 0
  Flags: none
  Next-hop type: reject               Index: 14       Reference: 1

Destination:  127.0.0.1/32
  Route type: interface
  Route reference: 0                  Route interface-index: 0
  Flags: sent to PFE
  Nexthop: 127.0.0.1
  Next-hop type: local                Index: 320      Reference: 1

...

Routing table: private1__.inet [Index 1]
Internet:

Destination:  default
  Route type: permanent
  Route reference: 0                  Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject               Index: 46       Reference: 1

Destination:  10.0.0.0/8
  Route type: interface
  Route reference: 0                  Route interface-index: 3
  Flags: sent to PFE
  Next-hop type: resolve              Index: 136      Reference: 1
  Next-hop interface: fxp1.0

...

Routing table: iso [Index 0]
ISO:

Destination:  default
  Route type: permanent
```

```
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                 Index: 38       Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination:  default
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                 Index: 22       Reference: 1

Destination:  ff00::/8
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: multicast discard      Index: 21       Reference: 1

...

Routing table: private1__.inet6 [Index 1]
Internet6:

Destination:  default
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                 Index: 54       Reference: 1

Destination:  fe80::2a0:a5ff:fe3d:375/128
  Route type: interface
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Nexthop: fe80::2a0:a5ff:fe3d:375
  Next-hop type: local                  Index: 75       Reference: 1

...
```

show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```
so-1/1/0 {
    unit 0 {
        family inet {
            rpf-check;
            address 192.0.2.2/30;
        }
    }
}

user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:
...
...
Destination:  192.0.2.3/32
  Route type: destination
  Route reference: 0                    Route interface-index: 67
```


Copyright © 2017, Juniper Networks, Inc.

```
  Flags: sent to PFE
  Nexthop: 192.0.2.3
  Next-hop type: broadcast             Index: 328      Reference: 1
  Next-hop interface: so-1/1/0.0
  RPF interface: so-1/1/0.0
```

show route forwarding-table family mpls

```
user@host> show route forwarding-table family mpls
Routing table: mpls
MPLS:
Destination        Type RtRef Next hop          Type Index NhRef Netif
default            perm     0                   rjct    19     1
0                  user     0                   recv    18     3
1                  user     0                   recv    18     3
2                  user     0                   recv    18     3
100000             user     0 10.31.1.6         swap  100001     fe-1/1/0.0
800002             user     0                   Pop              vt-0/3/0.32770

vt-0/3/0.32770 (VPLS)
                   user     0                   indr   351     4
                                                Push 800000, Push 100002(top)
so-0/0/0.0
```

show route forwarding-table family vpls

```
user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Destination        Type RtRef Next hop          Type Index NhRef Netif
default            dynm     0                   flood   353     1
default            perm     0                   rjct    298     1
fe-0/1/0.0         dynm     0                   flood   355     1
00:00:5E:00:53:1f/48             <<<<<Remote CE

                   dynm     0                   indr   351     4
                                                Push 800000, Push 100002(top)
so-0/0/0.0
00:00:5E:00:53:1f/48             <<<<<<Local CE

                   dynm     0                   ucst   354     2 fe-0/1/0.0
```

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled)

```
user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                   dscd      519     1
lsi.1048832        intf     0                   indr  1048574     4
                              172.16.3.2           Push 262145      621     2
ge-3/0/0.0
00:00:5E:00:53:01/48 user     0                  ucst      590     5 ge-2/3/9.0
0x30003/51         user     0                   comp      627     2
ge-2/3/9.0         intf     0                   ucst      590     5 ge-2/3/9.0
ge-3/1/3.0         intf     0                   ucst      619     4 ge-3/1/3.0
0x30002/51         user     0                   comp      600     2
0x30001/51         user     0                   comp      597     2
```

Copyright © 2017, Juniper Networks, Inc.

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics)

```
user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing, MAC Stats
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                    dscd      519     1
lsi.1048834        intf     0                    indr  1048574     4
                              172.16.3.2          Push 262145      592     2
ge-3/0/0.0
00:19:e2:25:d0:01/48 user     0                 ucst      590     5 ge-2/3/9.0
0x30003/51         user     0                    comp      630     2
ge-2/3/9.0         intf     0                    ucst      590     5 ge-2/3/9.0
ge-3/1/3.0         intf     0                    ucst      591     4 ge-3/1/3.0
0x30002/51         user     0                    comp      627     2
0x30001/51         user     0                    comp      624     2
```

show route forwarding-table family vpls extensive

```
user@host> show route forwarding-table family vpls extensive
Routing table: green.vpls [Index 2]
VPLS:

Destination:  default
  Route type: dynamic
  Route reference: 0                    Route interface-index: 72
  Flags: sent to PFE
  Next-hop type: flood                  Index: 289      Reference: 1
  Next-hop type: unicast                Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0
  Next-hop type: unicast                Index: 290      Reference: 3
  Next-hop interface: fe-0/1/2.0

Destination:  default
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: none
  Next-hop type: discard                Index: 341      Reference: 1

Destination:  fe-0/1/2.0
  Route type: dynamic
  Route reference: 0                    Route interface-index: 69
  Flags: sent to PFE
  Next-hop type: flood                  Index: 293      Reference: 1
  Next-hop type: indirect               Index: 363      Reference: 4
  Next-hop type: Push 800016
  Next-hop interface: at-1/0/1.0
  Next-hop type: indirect               Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
  Next-hop type: unicast                Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0

Destination:  fe-0/1/3.0
  Route type: dynamic
  Route reference: 0                    Route interface-index: 70
  Flags: sent to PFE
  Next-hop type: flood                  Index: 292      Reference: 1
```


Copyright © 2017, Juniper Networks, Inc.

```
  Next-hop type: indirect              Index: 363      Reference: 4
  Next-hop type: Push 800016
  Next-hop interface: at-1/0/1.0
  Next-hop type: indirect              Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
  Next-hop type: unicast               Index: 290      Reference: 3
  Next-hop interface: fe-0/1/2.0

Destination:  00:00:5E:00:53:01/48
  Route type: dynamic
  Route reference: 0                   Route interface-index: 70
  Flags: sent to PFE, prefix load balance
  Next-hop type: unicast               Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0
  Route used as destination:
    Packet count:         6640    Byte count:      675786
  Route used as source
    Packet count:         6894    Byte count:      696424

Destination:  00:00:5E:00:53:04/48
  Route type: dynamic
  Route reference: 0                   Route interface-index: 69
  Flags: sent to PFE, prefix load balance
  Next-hop type: unicast               Index: 290      Reference: 3
  Next-hop interface: fe-0/1/2.0
  Route used as destination:
    Packet count:           96    Byte count:        8079
  Route used as source:
    Packet count:          296    Byte count:       24955

Destination:  00:00:5E:00:53:05/48
  Route type: dynamic
  Route reference: 0                   Route interface-index: 74
  Flags: sent to PFE, prefix load balance
  Next-hop type: indirect              Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
```

show route forwarding-table table default

```
user@host> show route forwarding-table table default
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    36     2
0.0.0.0/32         perm     0                    dscd    34     1
10.0.60.0/30       user     0 10.0.60.13         ucst   713     5 fe-0/1/3.0
10.0.60.12/30      intf     0                    rslv   688     1 fe-0/1/3.0
10.0.60.12/32      dest     0 10.0.60.12         recv   686     1 fe-0/1/3.0
10.0.60.13/32      dest     0 0:5:85:8b:bc:22    ucst   713     5 fe-0/1/3.0
10.0.60.14/32      intf     0 10.0.60.14         locl   687     2
10.0.60.14/32      dest     0 10.0.60.14         locl   687     2
10.0.60.15/32      dest     0 10.0.60.15         bcst   685     1 fe-0/1/3.0
10.0.67.12/30      user     0 10.0.60.13         ucst   713     5 fe-0/1/3.0
10.0.80.0/30       ifdn     0 ff.3.0.21          ucst   676     1 so-0/0/1.0
10.0.80.0/32       dest     0 10.0.80.0          recv   678     1 so-0/0/1.0
10.0.80.2/32       user     0                    rjct    36     2
10.0.80.2/32       intf     0 10.0.80.2          locl   675     1
```

Copyright © 2017, Juniper Networks, Inc.

```
10.0.80.3/32       dest     0 10.0.80.3          bcst   677     1 so-0/0/1.0
10.0.90.12/30      intf     0                    rslv   684     1 fe-0/1/0.0
10.0.90.12/32      dest     0 10.0.90.12         recv   682     1 fe-0/1/0.0
10.0.90.14/32      intf     0 10.0.90.14         locl   683     2
10.0.90.14/32      dest     0 10.0.90.14         locl   683     2
10.0.90.15/32      dest     0 10.0.90.15         bcst   681     1 fe-0/1/0.0
10.5.0.0/16        user     0 192.168.187.126    ucst   324    15 fxp0.0
10.10.0.0/16       user     0 192.168.187.126    ucst   324    15 fxp0.0
10.13.10.0/23      user     0 192.168.187.126    ucst   324    15 fxp0.0
10.84.0.0/16       user     0 192.168.187.126    ucst   324    15 fxp0.0
10.150.0.0/16      user     0 192.168.187.126    ucst   324    15 fxp0.0
10.157.64.0/19     user     0 192.168.187.126    ucst   324    15 fxp0.0
10.209.0.0/16      user     0 192.168.187.126    ucst   324    15 fxp0.0

...

Routing table: default.iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    60     1

Routing table: default.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    44     1
::/128             perm     0                    dscd    42     1
ff00::/8           perm     0                    mdsc    43     1
ff02::1/128        perm     0 ff02::1            mcst    39     1

Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd    50     1
```

show route forwarding-table table logical-system-name/routing-instance-name

```
user@host> show route forwarding-table table R4/vpn-red
Logical system: R4
Routing table: vpn-red.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   563     1
0.0.0.0/32         perm     0                    dscd   561     2
172.16.0.1/32         user     0                    dscd   561     2
172.16.2.0/24         intf     0                    rslv   771     1 ge-1/2/0.3
172.16.2.0/32        dest     0 172.16.2.0           recv   769     1 ge-1/2/0.3
172.16.2.1/32         intf     0 172.16.2.1           locl   770     2
172.16.2.1/32         dest     0 172.16.2.1           locl   770     2
172.16.2.2/32         dest     0 0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0
                                                    ucst   789     1 ge-1/2/0.3
172.16.2.255/32       dest     0 172.16.2.255         bcst   768     1 ge-1/2/0.3
172.16.233.0/4         perm     1                    mdsc   562     1
172.16.233.1/32        perm     0 172.16.233.1         mcst   558     1
255.255.255.255/32 perm     0                    bcst   559     1

Logical system: R4
Routing table: vpn-red.iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   608     1
```


Copyright © 2017, Juniper Networks, Inc.

```
Logical system: R4
Routing table: vpn-red.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   708     1
::/128             perm     0                    dscd   706     1
ff00::/8           perm     0                    mdsc   707     1
ff02::1/128        perm     0 ff02::1            mcst   704     1

Logical system: R4
Routing table: vpn-red.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd   638
```

show route forwarding-table vpn

```
user@host> show route forwarding-table vpn VPN-A
      Routing table:: VPN-A.inet
Internet:
      Destination        Type RtRef Nexthop           Type Index NhRef Netif
      default            perm     0                   rjct     4     4
      10.39.10.20/30     intf     0 ff.3.0.21         ucst    40     1
      so-0/0/0.0
      10.39.10.21/32     intf     0 10.39.10.21       locl    36     1
      10.255.14.172/32   user     0                   ucst    69     2
      so-0/0/0.0
      10.255.14.175/32   user     0                   indr    81     3
                                                      Push 100004, Push
      100004(top) so-1/0/0.0
      172.16.233.0/4        perm     2                   mdsc     5     3
      172.16.233.1/32       perm     0 172.16.233.1         mcst     1     8
      172.16.233.5/32       user     1 172.16.233.5         mcst     1     8
      255.255.255.255/32 perm     0                   bcst     2     3
```

On QFX5200, the results for this command look like this:

```
show route forwarding-table family mpls
Routing table: default.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 65 1
0 user 0 recv 64 4
1 user 0 recv 64 4
2 user 0 recv 64 4
13 user 0 recv 64 4
300384 user 0 9.1.1.1 Pop 1711 2 xe-0/0/34.0
300384(S=0) user 0 9.1.1.1 Pop 1712 2 xe-0/0/34.0
300400 user 0 ulst 131071 2
                              10.1.1.2 Pop 1713 1 xe-0/0/38.0
                              172.16.11.2 Pop 1714 1 xe-0/0/40.0
300400(S=0) user 0 ulst 131072 2
                              10.1.1.2 Pop 1715 1 xe-0/0/38.0
                              172.16.11.2 Pop 1716 1 xe-0/0/40.0

Routing table: __mpls-oam__.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 1681 1
```

Copyright © 2017, Juniper Networks, Inc.

show route inactive-path

List of Syntax
Syntax on page 282
Syntax (EX Series Switches) on page 282

Syntax
show route inactive-path
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route inactive-path
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.

Options
none—Display all inactive routes.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show route active-path on page 200

List of Sample Output
show route inactive-path on page 282
show route inactive-path detail on page 283
show route inactive-path extensive on page 284
show route inactive-path terse on page 284

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route inactive-path

```
user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30     [OSPF/10] 03:57:28, metric 1
                    > via so-0/3/0.0
```


Copyright © 2017, Juniper Networks, Inc.

```
private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/8          [Direct/0] 04:39:56
                    > via fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30       [BGP/170] 04:38:17, localpref 100
                      AS path: 100 I
                    > to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route inactive-path detail

```
user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:     1
                Age: 3:58:24    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp1.0, selected
                State: <NotBest Int>
                Inactive reason: No difference
                Age: 4:40:52
                Task: IF
                AS path: I
```

```
red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)
         BGP    Preference: 170/-101
                Next-hop reference count: 6
                Source: 10.12.80.1
                Next hop: 10.12.80.1 via ge-6/3/2.0, selected
                State: <Ext>
                Inactive reason: Route Preference
                Peer AS:   100
                Age: 4:39:13
                Task: BGP_100.10.12.80.1+179
                AS path: 100 I
                Localpref: 100
                Router ID: 10.0.0.0
```

show route inactive-path extensive

The output for the **show route inactive-path extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see show route inactive-path detail on page 283.

show route inactive-path terse

```
user@host> show route inactive-path terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.12.100.12/30    O  10          1            >so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.0.0.0/8         D   0                     >fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.12.80.0/30      B 170        100            >10.12.80.1      100 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
```


Copyright © 2017, Juniper Networks, Inc.

```
private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

show route inactive-prefix

List of Syntax
Syntax on page 286
Syntax (EX Series Switches) on page 286

Syntax
show route inactive-prefix
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route inactive-prefix
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display inactive route destinations in each routing table.

Options
none—Display all inactive route destination.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route inactive-prefix on page 286
show route inactive-prefix detail on page 286
show route inactive-prefix extensive on page 287
show route inactive-prefix terse on page 287

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route inactive-prefix

```
user@host> show route inactive-prefix

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32        [Direct/0] 00:04:54
                    > via lo0.0
```

show route inactive-prefix detail

```
user@host> show route inactive-prefix detail

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

```
127.0.0.1/32 (1 entry, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Hidden Martian Int>
                Age: 4:51
                Task: IF
                AS path: I00:04:54
                    > via lo0.0
```

show route inactive-prefix extensive

The output for the **show route inactive-prefix extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see show route inactive-prefix detail on page 286.

show route inactive-prefix terse

```
user@host> show route inactive-prefix terse

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  127.0.0.1/32       D   0                       >lo0.0
```

Copyright © 2017, Juniper Networks, Inc.

show route instance

List of Syntax
Syntax on page 288
Syntax (EX Series Switches and QFX Series) on page 288

Syntax
show route instance
<brief | detail | summary>
<*instance-name*>
<logical-system (all | *logical-system-name*)>
<operational>

Syntax (EX Series Switches and QFX Series)
show route instance
<brief | detail | summary>
<*instance-name*>
<operational>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Display routing instance information.

Options
none—(Same as **brief**) Display standard information about all routing instances.

brief | detail | summary—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**. (These options are not available with the **operational** keyword.)

instance-name—(Optional) Display information for all routing instances whose name begins with this string (for example, **cust1**, **cust11**, and **cust111** are all displayed when you run the **show route instance cust1** command).

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

operational—(Optional) Display operational routing instances.

Required Privilege Level
view

Related Documentation
- *Example: Transporting IPv6 Traffic Across IPv4 Using Filter-Based Tunneling*
- *Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart*

List of Sample Output
show route instance on page 290
show route instance detail (Graceful Restart Complete) on page 290
show route instance detail (Graceful Restart Incomplete) on page 292
show route instance detail (VPLS Routing Instance) on page 293
show route instance operational on page 294
show route instance summary on page 294

Copyright © 2017, Juniper Networks, Inc.

Output Fields Table 27 on page 289 lists the output fields for the **show route instance** command. Output fields are listed in the approximate order in which they appear.

Table 27: show route instance Output Fields

Field Name	Field Description	Level of Output
Instance or *instance-name*	Name of the routing instance.	All levels
Operational Routing Instances	(**operational** keyword only) Names of all operational routing instances.	—
Type	Type of routing instance: **forwarding**, **l2vpn**, **no-forwarding**, **vpls**, **virtual-router**, or **vrf**.	All levels
State	State of the routing instance: **active** or **inactive**.	**brief detail** none
Interfaces	Name of interfaces belonging to this routing instance.	**brief detail** none
Restart State	Status of graceful restart for this instance: **Pending** or **Complete**.	**detail**
Path selection timeout	Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is **300**.	**detail**
Tables	Tables (and number of routes) associated with this routing instance.	**brief detail** none
Route-distinguisher	Unique route distinguisher associated with this routing instance.	**detail**
Vrf-import	VPN routing and forwarding instance import policy name.	**detail**
Vrf-export	VPN routing and forwarding instance export policy name.	**detail**
Vrf-import-target	VPN routing and forwarding instance import target community name.	**detail**
Vrf-export-target	VPN routing and forwarding instance export target community name.	**detail**
Vrf-edge-protection-id	Context identifier configured for edge-protection.	**detail**
Fast-reroute-priority	Fast reroute priority setting for a VPLS routing instance: **high**, **medium**, or **low**. The default is **low**.	**detail**
Restart State	Restart state: • **Pending:*protocol-name***—List of protocols that have not yet completed graceful restart for this routing table. • **Complete**—All protocols have restarted for this routing table.	**detail**
Primary rib	Primary table for this routing instance.	**brief** none **summary**
Active/holddown/hidden	Number of active, hold-down, and hidden routes.	All levels

Copyright © 2017, Juniper Networks, Inc.

Sample Output

show route instance

```
user@host> show route instance
Instance             Type
         Primary RIB                                    Active/holddown/hidden
master               forwarding
         inet.0                                         16/0/1
         iso.0                                          1/0/0
         mpls.0                                         0/0/0
         inet6.0                                        2/0/0
         l2circuit.0                                    0/0/0
__juniper_private1__ forwarding
         __juniper_private1__.inet.0                    12/0/0
         __juniper_private1__.inet6.0                   1/0/0
```

show route instance detail (Graceful Restart Complete)

```
user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding        State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0                : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0                 : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0           : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0               : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0           : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0       : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
```


Copyright © 2017, Juniper Networks, Inc.

```
    BGP-L.inet.0           : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
    BGP-L.mpls.0           : 3 routes (3 active, 0 holddown, 0 hidden)
    Restart Complete
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn             State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
LDP:
  Router ID: 10.69.105.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0             : 5 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
OSPF:
  Router ID: 10.69.101.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Vrf-import-target: [ target:11111
  Tables:
    OSPF.inet.0            : 8 routes (7 active, 0 holddown, 0 hidden)
    Restart Complete
RIP:
  Router ID: 10.69.102.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0             : 6 routes (6 active, 0 holddown, 0 hidden)
    Restart Complete
STATIC:
  Router ID: 10.69.100.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
```

Copyright © 2017, Juniper Networks, Inc.

```
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
```

show route instance detail (Graceful Restart Incomplete)

```
user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding        State: Active
  Restart State: Pending  Path selection timeout: 300
  Tables:
    inet.0                 : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0                  : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                 : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
    bgp.l3vpn.0            : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0            : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0        : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0           : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
    BGP-L.mpls.0           : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn             State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
```


Copyright © 2017, Juniper Networks, Inc.

```
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN L2VPN
LDP:
  Router ID: 10.69.105.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0             : 5 routes (4 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF LDP VPN
OSPF:
  Router ID: 10.69.101.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0            : 8 routes (7 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF VPN
RIP:
  Router ID: 10.69.102.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0             : 8 routes (6 active, 2 holddown, 0 hidden)
    Restart Pending: RIP VPN
STATIC:
  Router ID: 10.69.100.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
```

show route instance detail (VPLS Routing Instance)

```
user@host> show route instance detail test-vpls
test-vpls:
  Router ID: 0.0.0.0
  Type: vpls              State: Active
  Interfaces:
    lsi.1048833
```

```
    lsi.1048832
    fe-0/1/0.513
  Route-distinguisher: 10.255.37.65:1
  Vrf-import: [ __vrf-import-test-vpls-internal__ ]
  Vrf-export: [ __vrf-export-test-vpls-internal__ ]
  Vrf-import-target: [ target:300:1 ]
  Vrf-export-target: [ target:300:1 ]
  Vrf-edge-protection-id:  166.1.3.1  Fast-reroute-priority: high
  Tables:
    test-vpls.l2vpn.0          : 3 routes (3 active, 0 holddown, 0 hidden)
```

show route instance operational

```
user@host> show route instance operational
Operational Routing Instances:

master
default
```

show route instance summary

```
user@host> show route instance summary
Instance           Type           Primary rib            Active/holddown/hidden
master             forwarding
                                  inet.0                 15/0/1
                                  iso.0                  1/0/0
                                  mpls.0                 35/0/0
                                  l3vpn.0                0/0/0
                                  inet6.0                2/0/0
                                  l2vpn.0                0/0/0
                                  l2circuit.0            0/0/0
BGP-INET           vrf
                                  BGP-INET.inet.0        5/0/0
                                  BGP-INET.iso.0         0/0/0
                                  BGP-INET.inet6.0       0/0/0
BGP-L              vrf
                                  BGP-L.inet.0           5/0/0
                                  BGP-L.iso.0            0/0/0
                                  BGP-L.mpls.0           4/0/0
                                  BGP-L.inet6.0          0/0/0
L2VPN              l2vpn
                                  L2VPN.inet.0           0/0/0
                                  L2VPN.iso.0            0/0/0
                                  L2VPN.inet6.0          0/0/0
                                  L2VPN.l2vpn.0          2/0/0
LDP                vrf
                                  LDP.inet.0             4/0/0
                                  LDP.iso.0              0/0/0
                                  LDP.mpls.0             0/0/0
                                  LDP.inet6.0            0/0/0
                                  LDP.l2circuit.0        0/0/0
OSPF               vrf
                                  OSPF.inet.0            7/0/0
                                  OSPF.iso.0             0/0/0
                                  OSPF.inet6.0           0/0/0
RIP                vrf
                                  RIP.inet.0             6/0/0
                                  RIP.iso.0              0/0/0
                                  RIP.inet6.0            0/0/0
STATIC             vrf
                                  STATIC.inet.0          4/0/0
```


Copyright © 2017, Juniper Networks, Inc.

```
STATIC.iso.0            0/0/0
STATIC.inet6.0          0/0/0
```

Copyright © 2017, Juniper Networks, Inc.

show route label

List of Syntax
Syntax on page 296
Syntax (EX Series Switches) on page 296

Syntax
show route label *label*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route label *label*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.

Description
Display the routes based on a specified Multiprotocol Label Switching (MPLS) label value.

Options
label—Value of the MPLS label.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- *Example: Configuring Multipoint LDP In-Band Signaling for Point-to-Multipoint LSPs*

List of Sample Output
show route label terse on page 296
show route label on page 297
show route label detail on page 297
show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 297
show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 298
show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 298
show route label extensive on page 299

Output Fields
For information about output field table for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route label terse

```
user@host> show route label 100016 terse
```


Copyright © 2017, Juniper Networks, Inc.

```
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 100016             V 170                      >10.12.80.1
```

show route label

```
user@host> show route label 100016

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
100016             *[VPN/170] 03:25:41
                    > to 10.12.80.1 via ge-6/3/2.0, Pop
```

show route label detail

```
user@host> show route label 100016 detail

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
100016 (1 entry, 1 announced)
        *VPN    Preference: 170
                Next-hop reference count: 2
                Source: 10.12.80.1
                Next hop: 10.12.80.1 via ge-6/3/2.0, selected
                Label operation: Pop
                State: <Active Int Ext>
                Local AS:     1
                Age: 3:23:31
                Task: BGP.0.0.0.0+179
                Announcement bits (1): 0-KRT
                AS path: 100 I
                Ref Cnt: 2
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
                Next hop: via vt-0/1/0.1
                Next-hop index: 661
                Label operation: Pop
                Address: 0x9172130
                Next hop: via so-0/0/3.0
                Next-hop index: 654
                Label operation: Swap 299872
                State: **Active Int>
                Local AS:  1001
                Age: 8:20       Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
               FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
 src 192.168.142.2
```

Copyright © 2017, Juniper Networks, Inc.

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
                Next hop: via vt-0/1/0.1
                Next-hop index: 661
                Label operation: Pop
                Address: 0x9172130
                Next hop: via so-0/0/3.0
                Next-hop index: 654
                Label operation: Swap 299872
                State: **Active Int>
                Local AS:  1001
                Age: 8:20       Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
               FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
 src 192.168.142.2
```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Address: 0x2735208
                Next-hop reference count: 3
                Next hop type: Router, Next hop index: 1397
                Address: 0x2735d2c
                Next-hop reference count: 3
                Next hop: 1.3.8.2 via ge-1/2/22.0
                Label operation: Pop
                Load balance label:  None;
                Next hop type: Router, Next hop index: 1395
                Address: 0x2736290
                Next-hop reference count: 3
                Next hop: 1.3.4.2 via ge-1/2/18.0
                Label operation: Pop
                Load balance label:  None;
                State: <Active Int AckRequest MulticastRPF>
                Local AS:    10
                Age: 54:05      Metric: 1
                Validation State: unverified
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
               FECs bound to route: P2MP root-addr 1.1.1.1, grp: 232.1.1.1, src:
 192.168.219.11
                Primary Upstream : 1.1.1.3:0--1.1.1.2:0
                   RPF Nexthops  :
                        ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
                        ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
                Backup Upstream : 1.1.1.3:0--1.1.1.6:0
```


Copyright © 2017, Juniper Networks, Inc.

```
RPF Nexthops :
     ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xfffe
     ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xfffe
```

show route label extensive

The output for the show route label extensive command is identical to that of the **show route label detail** command. For sample output, see show route label detail on page 297.

Copyright © 2017, Juniper Networks, Inc.

show route label-switched-path

List of Syntax
Syntax on page 300
Syntax (EX Series Switches) on page 300

Syntax
show route label-switched-path *path-name*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route label-switched-path *path-name*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.5 for EX Series switches.

Description
Display the routes used in an MPLS label-switched path (LSP).

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output.

path-name—LSP tunnel name.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route label-switched-path on page 300

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route label-switched-path

```
user@host> show route label-switched-path sf-to-ny
inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          [MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
3.3.3.3/32         *[MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.2/32         *[MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
4.4.4.4/32         *[MPLS/7] 00:00:06, metric 0
                      to 111.222.1.9 via s0-0/0/0, label-switched-path abc
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path xyz
                      to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
```


Copyright © 2017, Juniper Networks, Inc.

```
111.222.1.9/32      [MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

Copyright © 2017, Juniper Networks, Inc.

show route martians

List of Syntax
Syntax on page 302
Syntax (EX Series Switches) on page 302

Syntax
show route martians
<logical-system (all | *logical-system-name*)>
<table *routing-table-name*>

Syntax (EX Series Switches)
show route martians
<table *routing-table-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the martian (invalid and ignored) entries associated with each routing table.

Options
none—Display standard information about route martians for all routing tables.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table *routing-table-name*—(Optional) Display information about route martians for all routing tables whose name begins with this string (for example, **inet.0** and **inet6.0** are both displayed when you run the **show route martians table inet** command).

Required Privilege Level
view

Related Documentation
- *Example: Configuring Class E Martian Addresses for Routing*
- *Understanding Martian Addresses*

List of Sample Output
show route martians on page 303

Output Fields
Table 28 on page 302 lists the output fields for the **show route martians** command. Output fields are listed in the approximate order in which they appear

Table 28: show route martians Output Fields

Field Name	Field Description
table-name	Name of the route table in which the route martians reside.
destination-prefix	Route destination.
match value	Route match parameter.
status	Status of the route: **allowed** or **disallowed**.

Copyright © 2017, Juniper Networks, Inc.

Sample Output

show route martians

```
user@host> show route martians

inet.0:
                0.0.0.0/0 exact -- allowed
                0.0.0.0/8 orlonger -- disallowed
                127.0.0.0/8 orlonger -- disallowed
                192.0.0.0/24 orlonger -- disallowed
                240.0.0.0/4 orlonger -- disallowed
                224.0.0.0/4 exact -- disallowed
                224.0.0.0/24 exact -- disallowed

inet.1:
                0.0.0.0/0 exact -- allowed
                0.0.0.0/8 orlonger -- disallowed
                127.0.0.0/8 orlonger -- disallowed
                192.0.0.0/24 orlonger -- disallowed
                240.0.0.0/4 orlonger -- disallowed

inet.2:
                0.0.0.0/0 exact -- allowed
                0.0.0.0/8 orlonger -- disallowed
                127.0.0.0/8 orlonger -- disallowed
                192.0.0.0/24 orlonger -- disallowed
                240.0.0.0/4 orlonger -- disallowed
                224.0.0.0/4 exact -- disallowed
                224.0.0.0/24 exact -- disallowed

inet.3:
                0.0.0.0/0 exact -- allowed
                0.0.0.0/8 orlonger -- disallowed
                127.0.0.0/8 orlonger -- disallowed
                192.0.0.0/24 orlonger -- disallowed
                240.0.0.0/4 orlonger -- disallowed
                224.0.0.0/4 exact -- disallowed
                224.0.0.0/24 exact -- disallowed
...

inet6.0:
                ::1/128 exact -- disallowed
                ff00::/8 exact -- disallowed
                ff02::/16 exact -- disallowed

inet6.1:
                ::1/128 exact -- disallowed

inet6.2:
                ::1/128 exact -- disallowed
                ff00::/8 exact -- disallowed
                ff02::/16 exact -- disallowed

inet6.3:
                ::1/128 exact -- disallowed
                ff00::/8 exact -- disallowed
                ff02::/16 exact -- disallowed
...
```

Copyright © 2017, Juniper Networks, Inc.

show route next-hop

List of Syntax
Syntax on page 304
Syntax (EX Series Switches) on page 304

Syntax
show route next-hop *next-hop*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route next-hop *next-hop*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the entries in the routing table that are being sent to the specified next-hop address.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of ouput.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

next-hop—Next-hop address.

Required Privilege Level
view

List of Sample Output
show route next-hop on page 304
show route next-hop detail on page 305
show route next-hop extensive on page 307
show route next-hop terse on page 308

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route next-hop

```
user@host> show route next-hop 192.168.71.254

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
172.16.0.0/12      *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16     *[Static/5] 06:26:25
```


Copyright © 2017, Juniper Networks, Inc.

```
                    > to 192.168.71.254 via fxp0.0
192.168.102.0/23   *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
207.17.136.0/24    *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
207.17.136.192/32  *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop detail

```
user@host> show route next-hop 192.168.71.254 detail

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
```

Copyright © 2017, Juniper Networks, Inc.

```
                AS path: I

192.168.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

192.168.102.0/23 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

207.17.136.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

207.17.136.192/32 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

show route next-hop extensive

```
user@host> show route next-hop 192.168.71.254 extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

192.168.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

192.168.102.0/23 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.102.0/23 -> {192.168.71.254}
        *Static Preference: 5
```

Copyright © 2017, Juniper Networks, Inc.

```
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

207.17.136.192/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.192/32 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop terse

```
user@host> show route next-hop 192.168.71.254 terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                     >192.168.71.254
* 10.209.0.0/16      S   5                     >192.168.71.254
* 172.16.0.0/12      S   5                     >192.168.71.254
```


Copyright © 2017, Juniper Networks, Inc.

```
* 192.168.0.0/16     S   5                       >192.168.71.254
* 192.168.102.0/23   S   5                       >192.168.71.254
* 207.17.136.0/24    S   5                       >192.168.71.254
* 207.17.136.192/32  S   5                       >192.168.71.254

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

show route no-community

List of Syntax
Syntax on page 310
Syntax (EX Series Switches) on page 310

Syntax
show route no-community
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route no-community
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the route entries in each routing table that are not associated with any community.

Options
none—(Same as **brief**) Display the route entries in each routing table that are not associated with any community.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route no-community on page 310
show route no-community detail on page 311
show route no-community extensive on page 311
show route no-community terse on page 312

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route no-community

```
user@host> show route no-community
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 00:36:27
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 00:36:27
                    > to 192.168.71.254 via fxp0.0
10.255.71.52/32    *[Direct/0] 00:36:27
                    > via lo0.0
10.255.71.63/32    *[OSPF/10] 00:04:39, metric 1
                    > to 35.1.1.2 via ge-3/1/0.0
10.255.71.64/32    *[OSPF/10] 00:00:08, metric 2
```


Copyright © 2017, Juniper Networks, Inc.

```
                     > to 35.1.1.2 via ge-3/1/0.0
10.255.71.240/32   *[OSPF/10] 00:05:04, metric 2
                      via so-0/1/2.0
                     > via so-0/3/2.0
10.255.71.241/32   *[OSPF/10] 00:05:14, metric 1
                     > via so-0/1/2.0
10.255.71.242/32   *[OSPF/10] 00:05:19, metric 1
                     > via so-0/3/2.0
172.16.12.0/24        *[OSPF/10] 00:05:14, metric 2
                     > via so-0/3/2.0
172.16.14.0/24        *[OSPF/10] 00:00:08, metric 3
                     > to 35.1.1.2 via ge-3/1/0.0
                       via so-0/1/2.0
                       via so-0/3/2.0
172.16.16.0/24        *[OSPF/10] 00:05:14, metric 2
                     > via so-0/1/2.0
.....
```

show route no-community detail

```
user@host> show route no-community detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 38:08
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 38:08
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
....
```

show route no-community extensive

```
user@host> show route no-community extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:03:33
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```

Copyright © 2017, Juniper Networks, Inc.

```
10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:03:33
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```

show route no-community terse

```
user@host> show route no-community terse

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                       >192.168.71.254
* 10.209.0.0/16      S   5                       >192.168.71.254
* 10.255.71.52/32    D   0                       >lo0.0
* 10.255.71.63/32    O  10          1            >35.1.1.2
* 10.255.71.64/32    O  10          2            >35.1.1.2
* 10.255.71.240/32   O  10          2             so-0/1/2.0
                                                 >so-0/3/2.0
* 10.255.71.241/32   O  10          1            >so-0/1/2.0
* 10.255.71.242/32   O  10          1            >so-0/3/2.0
* 172.16.12.0/24        O  10          2            >so-0/3/2.0
* 172.16.14.0/24        O  10          3            >35.1.1.2
                                                  so-0/1/2.0
                                                  so-0/3/2.0
* 172.16.16.0/24        O  10          2            >so-0/1/2.0
...
```


Copyright © 2017, Juniper Networks, Inc.

show route protocol

List of Syntax
Syntax on page 313
Syntax (EX Series Switches) on page 313

Syntax
show route protocol *protocol*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route protocol *protocol*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
ospf2 and **ospf3** options introduced in Junos OS Release 9.2.
ospf2 and **ospf3** options introduced in Junos OS Release 9.2 for EX Series switches.
flow option introduced in Junos OS Release 10.0.
flow option introduced in Junos OS Release 10.0 for EX Series switches.

Description
Display the route entries in the routing table that were learned from a particular protocol.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

protocol—Protocol from which the route was learned:

- **access**—Access route for use by DHCP application
- **access-internal**—Access-internal route for use by DHCP application
- **aggregate**—Locally generated aggregate route
- **arp**—Route learned through the Address Resolution Protocol
- **atmvpn**—Asynchronous Transfer Mode virtual private network
- **bgp**—Border Gateway Protocol
- **ccc**—Circuit cross-connect
- **direct**—Directly connected route
- **dvmrp**—Distance Vector Multicast Routing Protocol
- **esis**—End System-to-Intermediate System
- **flow**—Locally defined flow-specification route
- **frr**—Precomputed protection route or backup route used when a link goes down
- **isis**—Intermediate System-to-Intermediate System
- **ldp**—Label Distribution Protocol
- **l2circuit**—Layer 2 circuit

Copyright © 2017, Juniper Networks, Inc.

- **l2vpn**—Layer 2 virtual private network
- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network

NOTE: EX Series switches run a subset of these protocols. See the switch CLI for details.

Required Privilege Level view

List of Sample Output

show route protocol access on page 315
show route protocol access-internal extensive on page 315
show route protocol arp on page 315
show route protocol bgp on page 316
show route protocol bgp detail on page 316
show route protocol bgp detail (Labeled Unicast) on page 316
show route protocol bgp extensive on page 317
show route protocol bgp terse on page 318
show route protocol direct on page 318
show route protocol frr on page 318
show route protocol l2circuit detail on page 319
show route protocol l2vpn extensive on page 320
show route protocol ldp on page 320
show route protocol ldp extensive on page 321
show route protocol ospf (Layer 3 VPN) on page 322
show route protocol ospf detail on page 322
show route protocol rip on page 323

Copyright © 2017, Juniper Networks, Inc.

show route protocol rip detail on page 323
show route protocol ripng table inet6 on page 323
show route protocol static detail on page 324

Output Fields For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route protocol access

```
user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
```

show route protocol access-internal extensive

```
user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
        *Access-internal Preference: 12
                Next-hop reference count: 200000
                Next hop: 13.160.0.2 via fe-0/0/0.0, selected
                State: <Active Int>
       Age: 36
                Task: RPD Unix Domain Server./var/run/rpd_serv.local
                Announcement bits (1): 0-KRT
                AS path: I
```

show route protocol arp

```
user@host> show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.4/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.5/32        [ARP/4294967293] 00:04:32, from 20.20.1.1
                      Unusable
20.20.1.6/32        [ARP/4294967293] 00:04:34, from 20.20.1.1
                      Unusable
20.20.1.7/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.8/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
```

Copyright © 2017, Juniper Networks, Inc.

```
                      Unusable
20.20.1.9/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.10/32       [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.11/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
20.20.1.12/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
20.20.1.13/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
...
```

show route protocol bgp

```
user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21       *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
                      AS path: 10458 14203 2914 4788 4788 I
                    > to 192.168.167.254 via fxp0.0
```

show route protocol bgp detail

```
user@host> show route protocol bgp 66.117.63.0/24 detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24  (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006436
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324
                Next hop: 192.168.167.254 via fxp0.0, selected
                Protocol next hop: 192.168.69.71
                Indirect next hop: 8e166c0 342
                State: <Active Ext>
                Local AS:    69 Peer AS: 10458
                Age: 6d 10:42:42        Metric2: 0
                Task: BGP_10458.192.168.69.71+179
                Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
                AS path: 10458 14203 2914 4788 4788 I
                Communities: 2914:410 2914:2403 2914:3400
                Accepted
                Localpref: 100
                Router ID: 207.17.136.192
```

show route protocol bgp detail (Labeled Unicast)

```
user@host> show route protocol bgp 1.1.1.8/32 detail
inet.0: 45 destinations, 46 routes (45 active, 0 holddown, 0 hidden)
1.1.1.8/32 (2 entries, 2 announced)
State:
*BGP Preference: 1/-101
Next hop type: Indirect, Next hop index: 0
Address: 0xc007f30
Next-hop reference count: 2
Source: 1.1.1.1
Next hop type: Router, Next hop index: 614
Next hop: 20.1.1.2 via ge-0/0/1.0, selected
Label-switched-path lsp1
```


Copyright © 2017, Juniper Networks, Inc.

```
Label operation: Push 1000126, Push 1000125, Push 1000124, Push 1000123, Push
299872(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 1000126: None; Label 1000125: None; Label 1000124: None;
 Label 1000123: None; Label 299872: None;
Label element ptr: 0xc007860
Label parent element ptr: 0xc0089a0
Label element references: 1
Label element child references: 0
Label element lsp id: 0
Session Id: 0x140
Protocol next hop: 1.1.1.4
Label operation: Push 1000126, Push 1000125, Push 1000124, Push 1000123(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl, prop-ttl
Load balance label: Label 1000126: None; Label 1000125: None; Label 1000124: None;
 Label 1000123: None;
Indirect next hop: 0xae8d300 1048576 INH Session ID: 0x142
State:
Local AS: 5 Peer AS: 5
Age: 22:43 Metric2: 2
Validation State: unverified
Task: BGP_5.1.1.1.1
Announcement bits (2): 0-KRT 7-Resolve tree 2
AS path: I
Accepted
Route Labels: 1000123(top) 1000124 1000125 1000126
Localpref: 100
Router ID: 1.1.1.1
```

show route protocol bgp extensive

```
user@host> show route protocol bgp 192.168.64.0/21 extensive

inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
    Nexthop: Self
    AS path: [69] 10458 14203 2914 4788 4788 I
    Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4.  Val: 1
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006502
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324
                Next hop: 192.168.167.254 via fxp0.0, selected
                Protocol next hop: 192.168.69.71
                Indirect next hop: 8e166c0 342
                State: <Active Ext>
                Local AS:    69 Peer AS: 10458
                Age: 6d 10:44:45        Metric2: 0
                Task: BGP_10458.192.168.69.71+179
                Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
                AS path: 10458 14203 2914 4788 4788 I
                Communities: 2914:410 2914:2403 2914:3400
                Accepted
                Localpref: 100
                Router ID: 207.17.136.192
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Indirect next hops: 1
                            Protocol next hop: 192.168.69.71
                            Indirect next hop: 8e166c0 342
                            Indirect path forwarding next hops: 1
                                    Next hop type: Router
                                    Next hop: 192.168.167.254 via fxp0.0
                            192.168.0.0/16 Originating RIB: inet.0
                              Node path count: 1
                              Forwarding nexthops: 1
                                    Nexthop: 192.168.167.254 via fxp0.0
```

show route protocol bgp terse

```
user@host> show route protocol bgp 192.168.64.0/21 terse

inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  192.168.64.0/21    B 170        100            >172.16.100.1    10023 21 I
```

show route protocol direct

```
user@host> show route protocol direct

inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.8.0/24         *[Direct/0] 17w0d 10:31:49
                       > via fe-1/3/1.0
10.255.165.1/32     *[Direct/0] 25w4d 04:13:18
                       > via lo0.0
172.16.30.0/24       *[Direct/0] 17w0d 23:06:26
                       > via fe-1/3/2.0
192.168.164.0/22    *[Direct/0] 25w4d 04:13:20
                       > via fxp0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
                    *[Direct/0] 25w4d 04:13:21
                       > via lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::10:255:165:1/128
                    *[Direct/0] 25w4d 04:13:21
                       > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
                    *[Direct/0] 25w4d 04:13:21
                       > via lo0.0
```

show route protocol frr

```
user@host> show route protocol frr
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

```
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.3 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.4 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32       *[FRR/200] 00:05:35, from 20.20.1.1
                    > to 20.20.1.5 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32       *[FRR/200] 00:05:37, from 20.20.1.1
                    > to 20.20.1.6 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.7 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.8 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.9 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32      *[FRR/200] 00:05:38, from 20.20.1.1
...
```

show route protocol l2circuit detail

```
user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via ge-2/0/0.0, selected
                Label operation: Pop       Offset: 4
                State: <Active Int>
                Local AS:    99
                Age: 9:52
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

ge-2/0/0.0 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000, Push 100000(top)[0] Offset: -4
                Protocol next hop: 10.245.255.63
                Push 100000 Offset: -4
                 Indirect next hop: 86af0c0 298
                State: <Active Int>
                Local AS:    99
                Age: 9:52
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Next hop: via so-1/1/2.0 weight 1, selected
                    Label-switched-path my-lsp
                    Label operation: Push 100000[0]
                    Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
                    State: <Active Int>
                    Local AS:    99
                    Age: 10:21
                    Task: l2 circuit
                    Announcement bits (1): 0-LDP
                    AS path: I
                    VC Label 100000, MTU 1500, VLAN ID 512
```

show route protocol l2vpn extensive

```
user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
        *L2VPN  Preference: 7
                Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
                Label operation: Pop       Offset: 4
                State: <Active Int>
                Local AS:    69
                Age: 7:48
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0.0      /16 -> {indirect(288)}
        *L2VPN  Preference: 7
                Next hop: via so-0/0/1.0, selected
                Label operation: Push 800000 Offset: -4
                Protocol next hop: 10.255.14.220
                Push 800000 Offset: -4
                 Indirect next hop: 85142a0 288
                State: <Active Int>
                Local AS:    69
                Age: 7:48
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:69:1 Layer2-info: encaps:PPP,
                control flags:2, mtu: 0
```

show route protocol ldp

```
user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```


Copyright © 2017, Juniper Networks, Inc.

```
192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0

private1__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064             *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Pop
100080             *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Swap 100000
```

show route protocol ldp extensive

```
user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
        State: <FlashAll>
        *LDP    Preference: 9
                Next-hop reference count: 3
                Next hop: via t1-4/0/0.0, selected
                Label operation: Push 100000
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
                AS path: I

192.168.17.1/32 (1 entry, 1 announced)
        State: <FlashAll>
        *LDP    Preference: 9
                Next-hop reference count: 3
                Next hop: via t1-4/0/0.0, selected
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
```

Copyright © 2017, Juniper Networks, Inc.

```
                AS path: I
                Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                Label operation: Pop
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I

100080 (1 entry, 1 announced)
TSI:
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                Label operation: Swap 100000
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 192.168.16.1/32
```

show route protocol ospf (Layer 3 VPN)

```
user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30       *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.39.1.8/30        [OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
10.255.14.171/32   *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.255.14.179/32   *[OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
172.16.233.5/32       *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30       [OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
10.255.14.173/32   *[OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
172.16.233.5/32       *[OSPF/10] 20:26:20, metric 1
```

show route protocol ospf detail

```
user@host> show route protocol ospf detail
```


Copyright © 2017, Juniper Networks, Inc.

```
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
         OSPF   Preference:  10
                Nexthop: via so-0/2/2.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Age: 6:25       Metric: 1
                Area: 0.0.0.0
                Task: VPN-AB-OSPF
                AS path: I
                Communities: Route-Type:0.0.0.0:1:0

...
```

show route protocol rip

```
user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32   *[RIP/100] 20:24:34, metric 2
                    > to 10.39.1.22 via t3-0/2/2.0
172.16.233.9/32       *[RIP/100] 00:03:59, metric 1
```

show route protocol rip detail

```
user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
        *RIP    Preference: 100
                Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
                State: <Active Int>
                Age: 20:25:02   Metric: 2
                Task: VPN-AB-RIPv2
                Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
                AS path: I
                Route learned from 10.39.1.22 expires in 96 seconds
```

show route protocol ripng table inet6

```
user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

Copyright © 2017, Juniper Networks, Inc.

```
1111::6/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

show route protocol static detail

```
user@host> show route protocol static detail
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.5.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.13.10.0/23 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```


Copyright © 2017, Juniper Networks, Inc.

show route range

List of Syntax
Syntax on page 325
Syntax (EX Series Switches) on page 325

Syntax
show route range
<brief | detail | extensive | terse>
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route range
<brief | detail | extensive | terse>
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display routing table entries using a prefix range.

Options
none—Display standard information about all routing table entries using a prefix range.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

destination-prefix—Destination and prefix mask for the range.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route range on page 325
show route range destination-prefix on page 326
show route range detail on page 326
show route range extensive on page 327
show route range terse on page 328

Output Fields
For information about output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route range

```
user@host> show route range

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 00:30:01
```

Copyright © 2017, Juniper Networks, Inc.

```
                    > to 192.168.71.254 via fxp0.0
10.255.71.14/32    *[Direct/0] 00:30:01
                    > via lo0.0
172.16.0.0/12      *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16     *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21    *[Direct/0] 00:30:01
                    > via fxp0.0
192.168.71.14/32   *[Local/0] 00:30:01
                      Local via fxp0.0
192.168.102.0/23   *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
...
```

show route range destination-prefix

```
user@host> show route range 192.168.0.0/16

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16     *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21    *[Direct/0] 00:31:14
                    > via fxp0.0
192.168.71.14/32   *[Local/0] 00:31:14
                      Local via fxp0.0
192.168.102.0/23   *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
```

show route range detail

```
user@host> show route range detail

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.255.71.14/32 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
```


Copyright © 2017, Juniper Networks, Inc.

```
                State: <Active Int>
                Age: 30:05
                Task: IF
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

...
```

show route range extensive

```
user@host> show route range extensive

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:17
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:17
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.255.71.14/32 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Age: 30:17
                Task: IF
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Next hop: 192.168.71.254 via fxp0.0, selected
                    State: <Active NoReadvrt Int Ext>
                    Age: 30:17
                    Task: RT
                    Announcement bits (1): 0-KRT
                    AS path: I

...
```

show route range terse

```
user@host> show route range terse

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                     >192.168.71.254
* 10.209.0.0/16      S   5                     >192.168.71.254
* 10.255.71.14/32    D   0                     >lo0.0
* 172.16.0.0/12      S   5                     >192.168.71.254
* 192.168.0.0/16     S   5                     >192.168.71.254
* 192.168.64.0/21    D   0                     >fxp0.0
* 192.168.71.14/32   L   0                      Local
* 192.168.102.0/23   S   5                     >192.168.71.254
* 207.17.136.0/24    S   5                     >192.168.71.254
* 207.17.136.192/32  S   5                     >192.168.71.254

__juniper_private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.0.0.0/8         D   0                     >fxp2.0
                     D   0                     >fxp1.0
* 10.0.0.4/32        L   0                      Local

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  47.0005.80ff.f800.0000.0108.0001.0102.5507.1014/152
*                    D   0                     >lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  abcd::10:255:71:14/128
*                    D   0                     >lo0.0
  fe80::280:42ff:fe11:226f/128
*                    D   0                     >lo0.0

__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  fe80::280:42ff:fe11:226f/128
*                    D   0                     >lo0.16385
```


Copyright © 2017, Juniper Networks, Inc.

show route receive-protocol

List of Syntax
Syntax on page 329
Syntax (EX Series Switches) on page 329

Syntax
show route receive-protocol *protocol neighbor-address*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)

Syntax (EX Series Switches)
show route receive-protocol *protocol neighbor-address*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

protocol neighbor-address—Protocol transmitting the route (**bgp**, **dvmrp**, **msdp**, **pim**, **rip**, or **ripng**) and address of the neighboring router from which the route entry was received.

Additional Information
The output displays the selected routes and the attributes with which they were received, but does not show the effects of import policy on the routing attributes.

Required Privilege Level
view

List of Sample Output
show route receive-protocol bgp on page 332
show route receive-protocol bgp extensive on page 332
show route receive-protocol bgp table extensive on page 332
show route receive-protocol bgp logical-system extensive on page 333
show route receive-protocol bgp detail (Layer 2 VPN) on page 334
show route receive-protocol bgp extensive (Layer 2 VPN) on page 334
show route receive-protocol bgp (Layer 3 VPN) on page 335
show route receive-protocol bgp detail (Layer 3 VPN) on page 335
show route receive-protocol bgp detail (Long-Lived Graceful Restart) on page 336
show route receive-protocol bgp detail (Labeled Unicast) on page 336
show route receive-protocol bgp extensive (Layer 3 VPN) on page 337

Output Fields
Table 29 on page 330 describes the output fields for the **show route receive-protocol** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 29: show route receive-protocol Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table—for example, inet.0.	All levels
number destinations	Number of destinations for which there are routes in the routing table.	All levels
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** • **holddown** (routes that are in pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)	All levels
Prefix	Destination prefix.	none **brief**
MED	Multiple exit discriminator value included in the route.	none **brief**
destination-prefix (entry, announced)	Destination prefix. The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination.	**detail extensive**
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.	**detail extensive**
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table	**detail extensive**
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.	**detail extensive**
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.	**detail extensive**
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.	**detail extensive**
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routing devices by advertising VPN labels. VPN labels transit over either an RSVP or an LDP label-switched path (LSP) tunnel.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 29: show route receive-protocol Output Fields *(continued)*

Field Name	Field Description	Level of Output
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	All levels
Localpref or Lclpref	Local preference value included in the route.	All levels
AS path	Autonomous system (AS) path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the router, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.	All levels
Route Labels	Stack of labels carried in the BGP route update.	**detail extensive**
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.	**detail extensive**
Originator ID	(For route reflected output only) Address of routing device that originally sent the route to the route reflector.	**detail extensive**
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.	**detail extensive**
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.	**detail extensive**
Attrset AS	Number, local preference, and path of the AS that originated the route. These values are stored in the **Attrset** attribute at the originating routing device.	**detail extensive**
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	**detail extensive**
control flags	Control flags: **none** or **Site Down**.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 29: show route receive-protocol Output Fields *(continued)*

Field Name	Field Description	Level of Output
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	**detail extensive**

Sample Output

show route receive-protocol bgp

```
user@host> show route receive-protocol bgp 10.255.245.215

inet.0: 28 destinations, 33 routes (27 active, 0 holddown, 1 hidden)
Prefix                  Next hop              MED     Lclpref    AS path
10.22.1.0/24            10.255.245.215        0       100        I
10.22.2.0/24            10.255.245.215        0       100        I
```

show route receive-protocol bgp extensive

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
Prefix             Next hop               MED    Lclpref AS path
172.16.1.0/24 (1 entry, 1 announced)
     Next hop: 10.0.50.3
     Localpref: 100
     AS path: I <Orginator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.45
172.16.163.0/16 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.68
172.16.164.0/16 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.45
172.16.195.0/24 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix             Next hop               MED    Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Next hop               MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Next hop               MED    Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

show route receive-protocol bgp table extensive

```
user@host> show route receive-protocol bgp 207.17.136.192 table inet.0 66.117.68.0/24 extensive
inet.0: 227315 destinations, 227316 routes (227302 active, 0 holddown, 13 hidden)
* 66.117.63.0/24 (1 entry, 1 announced)
     Nexthop: 207.17.136.29
```


Copyright © 2017, Juniper Networks, Inc.

```
     Localpref: 100
     AS path: AS2 PA[6]: 14203 2914 3356 29748 33437 AS_TRANS
     AS path: AS4 PA[2]: 33437 393219
     AS path: Merged[6]: 14203 2914 3356 29748 33437 393219 I
     Communities: 2914:420
```

show route receive-protocol bgp logical-system extensive

```
user@host> show route receive-protocol bgp 10.0.0.9 logical-system PE4 extensive
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
* 10.0.0.0/30 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 10.0.0.4/30 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

  10.0.0.8/30 (2 entries, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 10.9.9.1/32 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 10.100.1.1/32 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 172.16.44.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300096
     Nexthop: 10.0.0.9
     AS path: 13979 I
     AIGP: 203

* 172.16.55.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300112
     Nexthop: 10.0.0.9
     AS path: 13979 7018 I
     AIGP: 25

* 172.16.66.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300144
     Nexthop: 10.0.0.9
     AS path: 13979 7018 I

* 172.16.99.0/24 (1 entry, 1 announced)
```

Copyright © 2017, Juniper Networks, Inc.

```
     Accepted
     Route Label: 300160
     Nexthop: 10.0.0.9
     AS path: 13979 7018 I
```

show route receive-protocol bgp detail (Layer 2 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171 detail
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix             Nexthop               MED    Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
```

show route receive-protocol bgp extensive (Layer 2 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171 extensive
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix             Nexthop               MED    Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
```


Copyright © 2017, Juniper Networks, Inc.

```
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
```

show route receive-protocol bgp (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171
      inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.175/32   10.255.14.171                     100 2 I
      10.255.14.179/32   10.255.14.171            2        100 I
      VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.175/32   10.255.14.171                     100 2 I
      10.255.14.177/32   10.255.14.171                     100 I
      iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.171:300:10.255.14.177/32
                         10.255.14.171                     100 I
      10.255.14.171:100:10.255.14.179/32
                         10.255.14.171            2        100 I
      10.255.14.171:200:10.255.14.175/32
                         10.255.14.171                     100 2 I
```

show route receive-protocol bgp detail (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.174 detail
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
     Route Distinguisher: 10.255.14.176:2
     VPN Label: 101264
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
         AS path: I
* 10.255.14.172/32 (1 entry, 1 announced)
     Route Distinguisher: 10.255.14.176:2
     VPN Label: 101280
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
```

Copyright © 2017, Juniper Networks, Inc.

```
        Localpref: 100
        AS path: I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 10.255.14.174:2:10.49.0.0/30 (1 entry, 0 announced)
     Route Distinguisher: 10.255.14.174:2
     VPN Label: 101264
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
         AS path: I
* 10.255.14.174:2:10.255.14.172/32 (1 entry, 0 announced)
     Route Distinguisher: 10.255.14.174:2
     VPN Label: 101280
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
         AS path: I
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

show route receive-protocol bgp detail (Long-Lived Graceful Restart)

```
user@host> show route receive-protocol bgp 10.4.12.11 detail

  bgp.l2vpn.0: 38 destinations, 39 routes (37 active, 0 holddown, 1 hidden)
  * 172.16.1.4:100:172.16.1.4/96 AD (1 entry, 1 announced)
    Accepted LongLivedStale LongLivedStaleImport
    Nexthop: 10.4.12.11
    Localpref: 100
    AS path: I
```

show route receive-protocol bgp detail (Labeled Unicast)

```
user@host> show route receive-protocol bgp 1.1.1.1 detail
inet.0: 45 destinations, 46 routes (45 active, 0 holddown, 0 hidden)
* 1.1.1.8/32 (2 entries, 2 announced)
Accepted
Route Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: 1.1.1.4
Localpref: 100
AS path: I
Entropy label capable, next hop field matches route next hop

inet.3: 15 destinations, 21 routes (6 active, 0 holddown, 14 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

inet6.0: 26 destinations, 28 routes (26 active, 0 holddown, 0 hidden)

* 100::1/128 (2 entries, 2 announced)
Accepted
Route Labels: 1000123(top) 1000124 1000125 1000126
```


Copyright © 2017, Juniper Networks, Inc.

```
Nexthop: ::ffff:1.1.1.4
Localpref: 100
AS path: I

inet6.3: 22 destinations, 23 routes (22 active, 0 holddown, 0 hidden)
```

show route receive-protocol bgp extensive (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     172.16.1.0/24 (1 entry, 1 announced)
          Nexthop: 10.0.50.3
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.45
     172.16.163.0/16 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.68
     172.16.164.0/16 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.45
     172.16.195.0/24 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.68
     inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

show route resolution

List of Syntax
Syntax on page 338
Syntax (EX Series Switches) on page 338

Syntax

show route resolution
<brief | detail | extensive | summary>
<index *index*>
<logical-system (all | *logical-system-name*)>
<*prefix*>
<table *routing-table-name*>
<unresolved>

Syntax (EX Series Switches)

show route resolution
<brief | detail | extensive | summary>
<index *index*>
<*prefix*>
<table *routing-table-name*>
<unresolved>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the entries in the next-hop resolution database. This database provides for recursive resolution of next hops through other prefixes in the routing table.

Options

none—Display standard information about all entries in the next-hop resolution database.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

index *index*—(Optional) Show the index of the resolution tree.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix network/destination-prefix—(Optional) Display database entries for the specified address.

table *routing-table-name*—(Optional) Display information about a particular routing table (for example, **inet.0**) where policy-based export is currently enabled.

unresolved—(Optional) Display routes that could not be resolved.

Required Privilege Level

view

Related Documentation

- *Example: Configuring Route Resolution on PE Routers*

List of Sample Output

show route resolution detail on page 339
show route resolution summary on page 340
show route resolution unresolved on page 340

Copyright © 2017, Juniper Networks, Inc.

Output Fields Table 30 on page 339 describes the output fields for the **show route resolution** command. Output fields are listed in the approximate order in which they appear.

Table 30: show route resolution Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table whose prefixes are resolved using the entries in the route resolution database. For routing table groups, this is the name of the primary routing table whose prefixes are resolved using the entries in the route resolution database.
Tree index	Tree index identifier.
Nodes	Number of nodes in the tree.
Reference count	Number of references made to the next hop.
Contributing routing tables	Routing tables used for next-hop resolution.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of **inet.0** resolving through **inet.0** and **inet.3**, this field indicates which routing table, **inet.0** or **inet.3**, provided the best path for a particular prefix.
Metric	Metric associated with the forwarding next hop.
Node path count	Number of nodes in the path.
Forwarding next hops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route resolution detail

```
user@host> show route resolution detail
Tree Index: 1, Nodes 0, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 2, Nodes 23, Reference Count 1
Contributing routing tables: inet.0 inet.3
10.10.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.0/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.1/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.1.4/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.5/32 Originating RIB: inet.0
```

Copyright © 2017, Juniper Networks, Inc.

```
  Node path count: 1
  Forwarding nexthops: 0
10.31.2.0/30 Originating RIB: inet.0
  Metric: 2 Node path count: 1
  Forwarding nexthops: 2
10.31.11.0/24 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
```

show route resolution summary

```
user@host> show route resolution summary
Tree Index: 1, Nodes 24, Reference Count 1
 Contributing routing tables: :voice.inet.0 :voice.inet.3
 Tree Index: 2, Nodes 2, Reference Count 1
 Contributing routing tables: inet.3
 Tree Index: 3, Nodes 43, Reference Count 1
 Contributing routing tables: inet.0 inet.3
```

show route resolution unresolved

```
user@host> show route resolution unresolved
Tree Index 1
vt-3/2/0.32769.0      /16
        Protocol Nexthop: 10.255.71.238 Push 800000
        Indirect nexthop: 0 -
vt-3/2/0.32772.0      /16
        Protocol Nexthop: 10.255.70.103 Push 800008
        Indirect nexthop: 0 -
Tree Index 2
```


Copyright © 2017, Juniper Networks, Inc.

show route snooping

Syntax

show route snooping
<brief | detail | extensive | terse>
<all>
<best *address/prefix*>
<exact *address*>
<logical-system *logical-system-name*>
<range *prefix-range*>
<summary>
<table *table-name*>

Release Information

Command introduced in Junos OS Release 8.5.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display the entries in the routing table that were learned from snooping.

Options

none—Display the entries in the routing table that were learned from snooping.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

all—(Optional) Display all entries, including hidden entries.

best *address/prefix*—(Optional) Display the longest match for the provided address and optional prefix.

exact *address/prefix*—(Optional) Display exact matches for the provided address and optional prefix.

logical-system *logical-system-name*—(Optional) Display information about a particular logical system, or type 'all'.

range *prefix-range*—(Optional) Display information for the provided address range.

summary—(Optional) Display route snooping summary statisitics.

table *table-name*—(Optional) Display information for the named table.

Required Privilege Level

view

List of Sample Output

show route snooping detail on page 341
show route snooping logical-system all on page 342

Output Fields

For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route snooping detail

```
user@host> show route snooping detail
```

```
__+domainAll__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

224.0.0.2/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next hop type: MultiRecv
                Next-hop reference count: 4
                State: <Active NoReadvrt Int>
                Age: 2:24
                Task: IGMP
                Announcement bits (1): 0-KRT
                AS path: I

224.0.0.22/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next hop type: MultiRecv
                Next-hop reference count: 4
                State: <Active NoReadvrt Int>
                Age: 2:24
                Task: IGMP
                Announcement bits (1): 0-KRT
                AS path: I

__+domainAll__.inet.1: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)

224.0.0.0.0.0.0.0/24 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4), Next hop index: 1048584
                Next-hop reference count: 4
                State: <Active Int>
                Age: 2:24
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

<snip>
```

show route snooping logical-system all

```
user@host> show route snooping logical-system all

logical-system: default

inet.1: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
Restart Unsupported
+ = Active Route, - = Last Active, * = Both

0.0,0.1,0.0,232.1.1.65,100.1.1.2/112*[Multicast/180] 00:07:36
                      Multicast (IPv4) Composite
0.0,0.1,0.0,232.1.1.66,100.1.1.2/112*[Multicast/180] 00:07:36
                      Multicast (IPv4) Composite
0.0,0.1,0.0,232.1.1.67,100.1.1.2/112*[Multicast/180] 00:07:36

<snip>

default-switch.inet.1: 237 dest, 237 rts (237 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

0.15,0.1,0.0,0.0.0.0,0.0.0.0,2/120*[Multicast/180] 00:08:21
                      Multicast (IPv4) Composite
0.15,0.1,0.0,0.0.0.0,0.0.0.0,2,17/128*[Multicast/180] 00:08:21
```


Copyright © 2017, Juniper Networks, Inc.

```
                                Multicast (IPv4) Composite

<snip>
```

Copyright © 2017, Juniper Networks, Inc.

show route source-gateway

List of Syntax
Syntax on page 344
Syntax (EX Series Switches) on page 344

Syntax
show route source-gateway *address*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route source-gateway *address*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the entries in the routing table that were learned from a particular address. The **Source** field in the **show route detail** command output lists the source for each route, if known.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

address—IP address of the system.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route source-gateway on page 344
show route source-gateway detail on page 345
show route source-gateway extensive on page 347

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route source-gateway

```
user@host> show route source-gateway 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
```


Copyright © 2017, Juniper Networks, Inc.

```
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
                   *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:2:3:1/96
                   *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
                   *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

10.255.70.103:2:3:1/96
                   *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3
```

show route source-gateway detail

```
user@host> show route source-gateway 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

```
Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-red-l2vpn
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down, mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS, control
flags:, mtu: 0
```


Copyright © 2017, Juniper Networks, Inc.

```
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Secondary Tables: green.l2vpn.0
10.255.70.103:2:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Secondary Tables: red.l2vpn.0
```

show route source-gateway extensive

```
user@host> show route source-gateway 10.255.70.103 extensive
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
```

Copyright © 2017, Juniper Networks, Inc.

```
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-red-l2vpn
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down, mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Secondary Tables: green.l2vpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.255.70.103 Metric: 2
                        Indirect next hop: 2 no-forward
                        Indirect path forwarding next hops: 1
                Next hop:        via so-0/3/0.0 weight 0x1
                        10.255.70.103/32 Originating RIB: inet.3
                          Metric: 2                     Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: via so-0/3/0.0

10.255.70.103:2:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-1
```


Copyright © 2017, Juniper Networks, Inc.

```
Route Distinguisher: 10.255.70.103:2
Next-hop reference count: 7
Source: 10.255.70.103
Protocol next hop: 10.255.70.103
Indirect next hop: 2 no-forward
State: <Active Int Ext>
Local AS:    69 Peer AS:    69
Age: 12:15:24   Metric2: 1
Task: BGP_69.10.255.70.103+179
AS path: I
Communities: target:11111:2 Layer2-info: encaps:VPLS,
control flags:Site-Down,
mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Secondary Tables: red.l2vpn.0
Indirect next hops: 1
        Protocol next hop: 10.255.70.103 Metric: 2
        Indirect next hop: 2 no-forward
        Indirect path forwarding next hops: 1
Next hop:        via so-0/3/0.0 weight 0x1
        10.255.70.103/32 Originating RIB: inet.3
          Metric: 2                      Node path count: 1
          Forwarding nexthops: 1
                Nexthop: via so-0/3/0.0
```

Copyright © 2017, Juniper Networks, Inc.

show route summary

List of Syntax
Syntax on page 350
Syntax (EX Series Switches) on page 350

Syntax
show route summary
<logical-system (all | *logical-system-name*)>
<table *routing-table-name*>

Syntax (EX Series Switches)
show route summary

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display summary statistics about the entries in the routing table.

CPU utilization might increase while the device learns routes. We recommend that you use the **show route summary** command after the device learns and enters the routes into the routing table. Depending on the size of your network, this might take several minutes. If you receive a "timeout communicating with routing daemon" error when using the **show route summary** command, wait several minutes before attempting to use the command again. This is not a critical system error, but you might experience a delay in using the command-line interface (CLI).

Options
none—Display summary statistics about the entries in the routing table.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table *routing-table-name*—(Optional) Display summary statistics for all routing tables whose name begins with this string (for example, **inet.0** and **inet6.0** are both displayed when you run the **show route summary table inet** command). If you only want to display statistics for a specific routing table, make sure to enter the exact name of that routing table.

Required Privilege Level
view

List of Sample Output
show route summary on page 352
show route summary table on page 352
show route summary table (with Route Limits Configured for the Routing Table) on page 353

Output Fields
Table 31 on page 350 lists the output fields for the **show route summary** command. Output fields are listed in the approximate order in which they appear.

Table 31: show route summary Output Fields

Field Name	Field Description
Router ID	Address of the local routing device.

Copyright © 2017, Juniper Networks, Inc.

Table 31: show route summary Output Fields *(continued)*

Field Name	Field Description
routing-table-name	Name of the routing table (for example, **inet.0**).
destinations	Number of destinations for which there are routes in the routing table.
routes	Number of routes in the routing table: • **active**—Number of routes that are active. • **holddown**—Number of routes that are in the hold-down state before being declared inactive. • **hidden**—Number of routes that are not used because of routing policy.

Restart complete	All protocols have restarted for this routing table.
	Restart state:
	• **Pending:*protocol-name***—List of protocols that have not yet completed graceful restart for this routing table.
	• **Complete**—All protocols have restarted for this routing table.
	For example, if the output shows-
	• `LDP.inet.0: 5 routes (4 active, 1 holddown, 0 hidden)`
	`Restart Pending: OSPF LDP VPN`
	This indicates that **OSPF**, **LDP**, and **VPN** protocols did not restart for **LDP.inet.0** routing table.
	• `vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)`
	`Restart Complete`
	This indicates that all protocols have restarted for **vpls_1.l2vpn.0** routing table.

Limit/Threshold	Displays the configured route limits for the routing table set with the maximum-prefixes and the maximum-paths statements. If you do not configure route limits for the routing table, the show output does not display this information. • **destinations**—The first number represents the maximum number of route prefixes installed in the routing table. The second number represents the number of route prefixes that trigger a warning message. • **routes**—The first number represents the maximum number of routes. The second number represents the number of routes that trigger a warning message.
Direct	Routes on the directly connected network.
Local	Local routes.
protocol-name	Name of the protocol from which the route was learned. For example, **OSPF**, **RSVP**, and **Static**.

Copyright © 2017, Juniper Networks, Inc.

Sample Output

show route summary

```
user@host> show route summary
Autonomous system number: 69
Router ID: 10.255.71.52
Maximum-ECMP: 32
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
              Direct:      6 routes,      5 active
               Local:      4 routes,      4 active
                OSPF:      5 routes,      4 active
              Static:      7 routes,      7 active
                IGMP:      1 routes,      1 active
                 PIM:      2 routes,      2 active

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
               RSVP:      2 routes,      2 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
              Direct:      1 routes,      1 active

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
               MPLS:      3 routes,      3 active
               VPLS:      4 routes,      2 active

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
              Direct:      2 routes,      2 active
                 PIM:      2 routes,      2 active
                 MLD:      1 routes,      1 active

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      2 routes,      2 active
              L2VPN:      2 routes,      2 active

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      2 routes,      2 active
              L2VPN:      1 routes,      1 active

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      4 routes,      4 active
```

show route summary table

```
user@host> show route summary table inet
Router ID: 192.168.0.1

inet.0: 32 destinations, 34 routes (31 active, 0 holddown, 1 hidden)
              Direct:      6 routes,      5 active
               Local:      9 routes,      9 active
                OSPF:      3 routes,      1 active
              Static:     13 routes,     13 active
                IGMP:      1 routes,      1 active
```


Copyright © 2017, Juniper Networks, Inc.

```
                  PIM:      2 routes,      2 active

inet.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
           Multicast:      1 routes,      1 active

inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
               Local:      1 routes,      1 active
                 PIM:      2 routes,      2 active

inet6.1: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
           Multicast:      1 routes,      1 active
```

show route summary table (with Route Limits Configured for the Routing Table)

```
user@host> show route summary table VPN-A.inet.0
Autonomous system number: 100
Router ID: 10.255.182.142

VPN-A.inet.0: 13 destinations, 14 routes (13 active, 0 holddown, 0 hidden)
Limit/Threshold: 2000/200 destinations 20/12 routes
              Direct:      2 routes,      2 active
               Local:      1 routes,      1 active
                OSPF:      4 routes,      3 active
                 BGP:      4 routes,      4 active
                IGMP:      1 routes,      1 active
                 PIM:      2 routes,      2 active
```

Copyright © 2017, Juniper Networks, Inc.

show route table

List of Syntax
Syntax on page 354
Syntax (EX Series Switches and QFX Series Switches) on page 354

Syntax

show route table *routing-table-name*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series Switches)

show route table *routing-table-name*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.
Show route table evpn statement introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.

Description

Display the route entries in a particular routing table.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

routing-table-name—Display route entries for all routing tables whose names begin with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route table inet** command).

Required Privilege Level

view

Related Documentation

- show route summary on page 350

List of Sample Output

show route table bgp.l2.vpn on page 365
show route table bgp.l3vpn.0 on page 365
show route table bgp.l3vpn.0 detail on page 365
show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 366
show route table bgp.evpn.0 on page 367
show route table evpna.evpn.0 on page 367
show route table inet.0 on page 367
show route table inet.3 on page 368
show route table inet.3 protocol ospf on page 368
show route table inet6.0 on page 368
show route table inet6.3 on page 369
show route table inetflow detail on page 369
show route table l2circuit.0 on page 369
show route table mpls on page 370

Copyright © 2017, Juniper Networks, Inc.

show route table mpls extensive on page 370
show route table mpls.0 on page 370
show route table mpls.0 detail (PTX Series) on page 371
show route table mpls.0 protocol ospf on page 372
show route table mpls.0 extensive (PTX Series) on page 372
show route table mpls.0 (RSVP Route—Transit LSP) on page 373
show route table vpls_1 detail on page 373
show route table vpn-a on page 373
show route table vpn-a.mdt.0 on page 374
show route table VPN-A detail on page 374
show route table VPN-AB.inet.0 on page 374
show route table VPN_blue.mvpn-inet6.0 on page 375
show route table vrf1.mvpn.0 extensive on page 375
show route table inetflow detail on page 376
show route table bgp.evpn.0 extensive |no-more (EVPN) on page 379

Output Fields Table 17 on page 193 describes the output fields for the **show route table** command. Output fields are listed in the approximate order in which they appear.

Table 32: show route table Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).

Restart complete	All protocols have restarted for this routing table.
	
	Restart state:
	
	• **Pending:*protocol-name***—List of protocols that have not yet completed graceful restart for this routing table.
	• **Complete**—All protocols have restarted for this routing table.
	
	For example, if the output shows-
	
	• `LDP.inet.0 : 5 routes (4 active, 1 holddown, 0 hidden)`
	`Restart Pending: OSPF LDP VPN`
	
	This indicates that **OSPF**, **LDP**, and **VPN** protocols did not restart for the **LDP.inet.0** routing table.
	• `vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)`
	`Restart Complete`
	
	This indicates that all protocols have restarted for the **vpls_1.l2vpn.0** routing table.

number destinations	Number of destinations for which there are routes in the routing table.
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)

Copyright © 2017, Juniper Networks, Inc.

Table 32: show route table Output Fields *(continued)*

Field Name	Field Description

route-destination (entry, announced)	Route destination (for example:10.0.0.1/24). The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• ***MPLS-label*** (for example, 80001).
	• ***interface-name*** (for example, ge-1/0/2).
	• ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.
	• ***inclusive multicast Ethernet tag route***—Type of route destination represented by (for example, 3:100.100.100.10:100::0::10::100.100.100.10/384):
	  ◦ *route distinguisher*—(8 octets) Route distinguisher (RD) must be the RD of the EVPN instance (EVI) that is advertising the NLRI.
	  ◦ *Ethernet tag ID*—(4 octets) Identifier of the Ethernet tag. Can set to 0 or to a valid Ethernet tag value.
	  ◦ *IP address length*—(1 octet) Length of IP address in bits.
	  ◦ *originating router's IP address*—(4 or 16 octets) Must set to the provider edge (PE) device's IP address. This address should be common for all EVIs on the PE device, and may be the PE device's loopback address.

label stacking	**(**Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.

Copyright © 2017, Juniper Networks, Inc.

Table 32: show route table Output Fields *(continued)*

Field Name	Field Description
Level	(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 20 on page 229.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 21 on page 231.

Copyright © 2017, Juniper Networks, Inc.

Table 32: show route table Output Fields *(continued)*

Field Name	Field Description
Local AS	AS number of the local routing devices.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric*n*	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.
Task	Name of the protocol that has added the route.
Announcement bits	The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the kernel routing table (KRT) for installing the route into the Packet Forwarding Engine, to a resolve tree, a Layer 2 VC, or even a VPN. For example, ***n*-Resolve inet** indicates that the specified route is used for route resolution for next hops found in the routing table. • ***n***—An index used by Juniper Networks customer support only.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **Recorded**—The AS path is recorded by the sample process (sampled). • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.

Copyright © 2017, Juniper Networks, Inc.

Table 32: show route table Output Fields *(continued)*

Field Name	Field Description
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Indicates point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, indicates the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, indicates the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 22 on page 233 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or **Site Down**.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).

Copyright © 2017, Juniper Networks, Inc.

Table 32: show route table Output Fields *(continued)*

Field Name	Field Description
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag.
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 20 on page 229 describes all possible values for the Next-hop Types output field.

Table 33: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.

Copyright © 2017, Juniper Networks, Inc.

Table 33: Next-hop Types Output Field Values *(continued)*

Next-Hop Type	Description
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrt)	Regular multicast next hop.
Router	A specific node or set of nodes to which the routing device forwards packets that match the route prefix. To qualify as a next-hop type router, the route must meet the following criteria: • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Copyright © 2017, Juniper Networks, Inc.

Table 21 on page 231 describes all possible values for the State output field. A route can be in more than one state (for example, **<Active NoReadvrt Int Ext>**).

Table 34: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGP path is available.

Copyright © 2017, Juniper Networks, Inc.

Table 34: State Output Field Values *(continued)*

Value	Description
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available.
Router ID	Path through a neighbor with lower ID is available.
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Copyright © 2017, Juniper Networks, Inc.

Table 22 on page 233 describes the possible values for the Communities output field.

Table 35: Communities Output Field Values

Value	Description
area-number	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is **0**. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: ***local AS number:link-bandwidth-number***	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
link-bandwidth-number	Link-bandwidth number: from **0** through **4,294,967,295** (bytes per second).
local AS number	Local AS number: from **1** through **65,535**.
options	1 byte. Currently this is only used if the route type is **5** or **7**. Setting the least significant bit in the field indicates that the route carries a type 2 metric.
origin	(Used with VPNs) Identifies where the route came from.
ospf-route-type	1 byte, encoded as **1** or **2** for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); **3** for summary routes; **5** for external routes (area number must be**0**); **7** for NSSA routes; or **129** for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x8000**. The format is ***area-number:ospf-route-type:options***.
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x0306**. The format is ***area-number:ospf-route-type:options***.
target	Defines which VPN the route participates in; **target** has the format ***32-bit IP address:16-bit number***. For example, 10.19.0.0:100.
unknown IANA	Incoming IANA codes with a value between **0x1** and **0x7fff**. This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above **0x8000**. This code of the BGP extended community attribute is accepted, but it is not recognized.

Copyright © 2017, Juniper Networks, Inc.

Sample Output

show route table bgp.l2.vpn

```
user@host> show route table bgp.l2.vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
                   *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
      bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.255.71.15:100:10.255.71.17/32
                         *[BGP/170] 00:03:59, MED 1, localpref 100, from
      10.255.71.15
                           AS path: I
                          > via so-2/1/0.0, Push 100020, Push 100011(top)
      10.255.71.15:200:10.255.71.18/32
                         *[BGP/170] 00:03:59, MED 1, localpref 100, from
      10.255.71.15
                           AS path: I
                          > via so-2/1/0.0, Push 100021, Push 100011(top)
```

show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:172.16.4.0/8 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
         Route Distinguisher: 10.255.245.12:1
         Source: 10.255.245.12
         Next hop: 192.168.208.66 via fe-0/0/0.0, selected
         Label operation: Push 182449
         Protocol next hop: 10.255.245.12
         Push 182449
         Indirect next hop: 863a630 297
         State: <Active Int Ext>
         Local AS:    35 Peer AS:    35
         Age: 12:19      Metric2: 1
         Task: BGP_35.10.255.245.12+179
         Announcement bits (1): 0-BGP.0.0.0.0+179
         AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

         Communities: 2914:420 target:11111:1 origin:56:78
         VPN Label: 182449
         Localpref: 100
          Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
         Route Distinguisher: 10.255.245.12:1
         Source: 10.255.245.12
         Next hop: 192.168.208.66 via fe-0/0/0.0, selected
```

Copyright © 2017, Juniper Networks, Inc.

```
            Label operation: Push 182465
            Protocol next hop: 10.255.245.12
            Push 182465
            Indirect next hop: 863a8f0 305
            State: <Active Int Ext>
            Local AS:    35 Peer AS:    35
            Age: 12:19      Metric2: 1
            Task: BGP_35.10.255.245.12+179
            Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
                   Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                   VPN Label: 182465
                   Localpref: 100
                   Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
            Route Distinguisher: 10.255.245.12:1
            Source: 10.255.245.12
            Next hop: 192.168.208.66 via fe-0/0/0.0, selected
            Label operation: Push 182465
            Protocol next hop: 10.255.245.12
            Push 182465
            Indirect next hop: 86bd210 330
            State: <Active Int Ext>
            Local AS:    35 Peer AS:    35
            Age: 12:19      Metric2: 1
            Task: BGP_35.10.255.245.12+179
            Announcement bits (1): 0-BGP.0.0.0.0+179
            AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

            6496 I
            Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
            VPN Label: 182465
            Localpref: 100
            Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
            Route Distinguisher: 10.255.245.12:1
            Source: 10.255.245.12
            Next hop: 192.168.208.66 via fe-0/0/0.0, selected
            Label operation: Push 182465
            Protocol next hop: 10.255.245.12
            Push 182465
            Indirect next hop: 86bd210 330
            State: <Active Int Ext>
            Local AS:    35 Peer AS:    35
            Age: 12:19      Metric2: 1
            Task: BGP_35.10.255.245.12+179
            Announcement bits (1): 0-BGP.0.0.0.0+179
            AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

            6496 I
            Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
            VPN Label: 182465
            Localpref: 100
```

show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```
user@host> show route table bgp.rtarget.o
```


Copyright © 2017, Juniper Networks, Inc.

```
bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100:100:100/96
                   *[RTarget/5] 00:03:14
                      Type Proxy
                        for 10.255.165.103
                        for 10.255.166.124
                      Local
```

show route table bgp.evpn.0

```
user@host> show route table bgp.evpn.0
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2:100.100.100.2:100::0::00:26:88:5f:67:b0/304
                   *[BGP/170] 11:00:05, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.2:100::0::00:51:51:51:51:51/304
                   *[BGP/170] 11:00:05, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.3:100::0::00:52:52:52:52:52/304
                   *[BGP/170] 10:59:58, localpref 100, from 100.100.100.3
                      AS path: I, validation-state: unverified
                    > to 100.64.13.3 via ge-2/0/8.0, label-switched-path R0toR2
2:100.100.100.3:100::0::a8:d0:e5:5b:01:c8/304
                   *[BGP/170] 10:59:58, localpref 100, from 100.100.100.3
                      AS path: I, validation-state: unverified
                    > to 100.64.13.3 via ge-2/0/8.0, label-switched-path R0toR2
3:100.100.100.2:100::1000::100.100.100.2/304
                   *[BGP/170] 11:00:16, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
3:100.100.100.2:100::2000::100.100.100.2/304
                   *[BGP/170] 11:00:16, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
```

show route table evpna.evpn.0

```
user@host> show route table evpna.evpn.0
evpna.evpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

3:100.100.100.10:100::0::10::100.100.100.10/384
                   *[EVPN/170] 01:37:09
                      Indirect
3:100.100.100.2:100::2000::100.100.100.2/304
                   *[EVPN/170] 01:37:12
                      Indirect
```

show route table inet.0

```
user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:51:57
                    > to 172.16.5.254 via fxp0.0
```

Copyright © 2017, Juniper Networks, Inc.

```
10.0.0.1/32          *[Direct/0] 00:51:58
                      > via at-5/3/0.0
10.0.0.2/32          *[Local/0] 00:51:58
                       Local
10.12.12.21/32      *[Local/0] 00:51:57
                       Reject
10.13.13.13/32      *[Direct/0] 00:51:58
                      > via t3-5/2/1.0
10.13.13.14/32      *[Local/0] 00:51:58
                       Local
10.13.13.21/32      *[Local/0] 00:51:58
                       Local
10.13.13.22/32      *[Direct/0] 00:33:59
                      > via t3-5/2/0.0
127.0.0.1/32         [Direct/0] 00:51:58
                      > via lo0.0
10.222.5.0/24       *[Direct/0] 00:51:58
                      > via fxp0.0
10.222.5.81/32      *[Local/0] 00:51:58
                       Local
```

show route table inet.3

```
user@host> show route table inet.3
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32          *[LDP/9] 00:25:43, metric 10, tag 200
                        to 10.2.94.2 via lt-1/2/0.49
                      > to 10.2.3.2 via lt-1/2/0.23
```

show route table inet.3 protocol ospf

```
user@host> show route table inet.3 protocol ospf
inet.3: 9 destinations, 18 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.20/32          [L-OSPF/10] 1d 00:00:56, metric 2
                     > to 10.0.10.70 via lt-1/2/0.14, Push 800020
                       to 10.0.6.60 via lt-1/2/0.12, Push 800020, Push 800030(top)
1.1.1.30/32          [L-OSPF/10] 1d 00:01:01, metric 3
                     > to 10.0.10.70 via lt-1/2/0.14, Push 800030
                       to 10.0.6.60 via lt-1/2/0.12, Push 800030
1.1.1.40/32          [L-OSPF/10] 1d 00:01:01, metric 4
                     > to 10.0.10.70 via lt-1/2/0.14, Push 800040
                       to 10.0.6.60 via lt-1/2/0.12, Push 800040
1.1.1.50/32          [L-OSPF/10] 1d 00:01:01, metric 5
                     > to 10.0.10.70 via lt-1/2/0.14, Push 800050
                       to 10.0.6.60 via lt-1/2/0.12, Push 800050
1.1.1.60/32          [L-OSPF/10] 1d 00:01:01, metric 6
                     > to 10.0.10.70 via lt-1/2/0.14, Push 800060
                       to 10.0.6.60 via lt-1/2/0.12, Pop
```

show route table inet6.0

```
user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0
```


Copyright © 2017, Juniper Networks, Inc.

```
fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0
```

show route table inet6.3

```
user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
                   *[LDP/9] 00:00:22, metric 1
                    > via so-1/0/0.0
::10.255.245.196/128
                   *[LDP/9] 00:00:08, metric 1
                    > via so-1/0/0.0, Push 100008
```

show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 2
                State: <Active Ext>
                Local AS: 64502 Peer AS: 64500
                Age: 4
                Task: BGP_64500.10.12.99.5+3792
                Announcement bits (1): 0-Flow
                AS path: 64500 I
                Communities: traffic-rate:0:0
                Validation state: Accept, Originator: 10.12.99.5
                Via: 10.12.44.0/24, Active
                Localpref: 100
                Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
        *Flow   Preference: 5
                Next-hop reference count: 2
                State: <Active>
                Local AS: 64502
                Age: 6:30
                Task: RT Flow
                Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
                AS path: I
                Communities: 1:1
```

show route table l2circuit.0

```
user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
```

Copyright © 2017, Juniper Networks, Inc.

```
10.1.1.195:CtrlWord:1:2:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
```

show route table mpls

```
user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1                  *[MPLS/0] 00:13:55, metric 1
                      Receive
2                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1024               *[VPN/0] 00:04:18
                      to table red.inet.0, Pop
```

show route table mpls extensive

```
user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
        *LDP    Preference:   9
                Next hop:  via so-1/0/0.0, selected
                Pop
                State: <Active Int>
                Age: 29:50      Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 10.0.0.194/32
```

show route table mpls.0

```
user@host> show route table mpls.0
mpls.0: 18 destinations, 19 routes (18 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 11:39:56, metric 1
                      to table inet.0
0(S=0)             *[MPLS/0] 11:39:56, metric 1
                      to table mpls.0
1                  *[MPLS/0] 11:39:56, metric 1
                      Receive
2                  *[MPLS/0] 11:39:56, metric 1
                      to table inet6.0
2(S=0)             *[MPLS/0] 11:39:56, metric 1
                      to table mpls.0
13                 *[MPLS/0] 11:39:56, metric 1
                      Receive
303168             *[EVPN/7] 11:00:49, routing-instance pbbn10, route-type
Ingress-MAC, ISID 0
                      to table pbbn10.evpn-mac.0
303184             *[EVPN/7] 11:00:53, routing-instance pbbn10, route-type
Ingress-IM, ISID 1000
```


Copyright © 2017, Juniper Networks, Inc.

```
                    to table pbbn10.evpn-mac.0
                   [EVPN/7] 11:00:53, routing-instance pbbn10, route-type
Ingress-IM, ISID 2000
                    to table pbbn10.evpn-mac.0
303264             *[EVPN/7] 11:00:53, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-IM, ISID 1000
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303280             *[EVPN/7] 11:00:53, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-IM, ISID 2000
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303328             *[EVPN/7] 11:00:49, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303344             *[EVPN/7] 11:00:49, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303360             *[EVPN/7] 11:00:47, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:26:88:5f:67:b0
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303376             *[EVPN/7] 11:00:47, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:51:51:51:51:51
                    > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303392             *[EVPN/7] 11:00:35, remote-pe 100.100.100.3, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303408             *[EVPN/7] 11:00:35, remote-pe 100.100.100.3, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303424             *[EVPN/7] 11:00:33, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC a8:d0:e5:5b:01:c8
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303440             *[EVPN/7] 11:00:33, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:52:52:52:52:52
                    > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
```

show route table mpls.0 detail (PTX Series)

```
user@host> show route table mpls.0 detail
ge-0/0/2.600 (1 entry, 1 announced)
        *L2VPN  Preference: 7
                Next hop type: Indirect
                Address: 0x9438f34
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 567
                Next hop: 10.0.0.1 via ge-0/0/1.0, selected
                Label operation: Push 299808
                Label TTL action: prop-ttl
                Load balance label: Label 299808:None;
                Session Id: 0x1
                Protocol next hop: 10.255.255.1
                Label operation: Push 299872 Offset: 252
                Label TTL action: no-prop-ttl
                Load balance label: Label 299872:Flow label PUSH;
                Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                State: <Active Int>
                Age: 21         Metric2: 1
                Validation State: unverified
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 2-Common L2 VC
                AS path: I
```

Copyright © 2017, Juniper Networks, Inc.

show route table mpls.0 protocol ospf

```
user@host> show route table mpls.0 protocol ospf
mpls.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299952             *[L-OSPF/10] 23:59:42, metric 0
                    > to 10.0.10.70 via lt-1/2/0.14, Pop
                      to 10.0.6.60 via lt-1/2/0.12, Swap 800070, Push 800030(top)
299952(S=0)        *[L-OSPF/10] 23:59:42, metric 0
                    > to 10.0.10.70 via lt-1/2/0.14, Pop
                      to 10.0.6.60 via lt-1/2/0.12, Swap 800070, Push 800030(top)
299968             *[L-OSPF/10] 23:59:48, metric 0
                    > to 10.0.6.60 via lt-1/2/0.12, Pop
```

show route table mpls.0 extensive (PTX Series)

```
user@host> show route table mpls.0 extensive
ge-0/0/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel ge-0/0/2.600.0      /32 -> {composite(570)}
        *L2VPN  Preference: 7
                Next hop type: Indirect
                Address: 0x9438f34
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 567
                Next hop: 10.0.0.1 via ge-0/0/1.0, selected
                Label operation: Push 299808
                Label TTL action: prop-ttl
                Load balance label: Label 299808:None;
                Session Id: 0x1
                Protocol next hop: 10.255.255.1
                Label operation: Push 299872 Offset: 252
                Label TTL action: no-prop-ttl
                Load balance label: Label 299872:Flow label PUSH;
                Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                State: <Active Int>
                Age: 47         Metric2: 1
                Validation State: unverified
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 2-Common L2 VC
                AS path: I
                Composite next hops: 1
                        Protocol next hop: 10.255.255.1 Metric: 1
                        Label operation: Push 299872 Offset: 252
                        Label TTL action: no-prop-ttl
                        Load balance label: Label 299872:Flow label PUSH;
                        Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                        Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.0.0.1 via ge-0/0/1.0
                                Session Id: 0x1
                        10.255.255.1/32 Originating RIB: inet.3
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 10.0.0.1 via ge-0/0/1.0
```


Copyright © 2017, Juniper Networks, Inc.

show route table mpls.0 (RSVP Route—Transit LSP)

```
user@host> show route table mpls.0

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:37:31, metric 1
                      Receive
1                  *[MPLS/0] 00:37:31, metric 1
                      Receive
2                  *[MPLS/0] 00:37:31, metric 1
                      Receive
13                 *[MPLS/0] 00:37:31, metric 1
                      Receive
300352             *[RSVP/7/1] 00:08:00, metric 1
                   > to 10.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0)        *[RSVP/7/1] 00:08:00, metric 1
                   > to 10.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384             *[RSVP/7/2] 00:05:20, metric 1
                    > to 10.64.1.106 via ge-1/0/0.0, Pop
300384(S=0)        *[RSVP/7/2] 00:05:20, metric 1
                    > to 10.64.1.106 via ge-1/0/0.0, Pop
```

show route table vpls_1 detail

```
user@host> show route table vpls_1 detail
vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

172.16.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF
```

show route table vpn-a

```
user@host> show route table vpn-a
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
                   *[VPN/7] 05:48:27
                     Discard
192.168.24.1:1:2:1/96
                   *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
                   *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

Copyright © 2017, Juniper Networks, Inc.

show route table vpn-a.mdt.0

```
user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
                   *[MVPN/70] 01:23:05, metric2 1
                      Indirect
1:1:1:10.255.14.218:232.1.1.1/144
                   *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
                      AS path: I
                    > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
                   *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
                      AS path: I
> via so-0/0/1.0, label-switched-path r0-to-r2
```

show route table VPN-A detail

```
user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.179.13:200
                Next hop type: Indirect
                Next-hop reference count: 5
                Source: 10.255.179.13
                Next hop type: Router, Next hop index: 732
                Next hop: 10.39.1.14 via fe-0/3/0.0, selected
                Label operation: Push 299824, Push 299824(top)
                Protocol next hop: 10.255.179.13
                Push 299824
                Indirect next hop: 8f275a0 1048574
                State: (Secondary Active Int Ext)
                Local AS: 1 Peer AS: 1
                Age: 3:41:06 Metric: 1 Metric2: 1
                Task: BGP_1.10.255.179.13+64309
                Announcement bits (2): 0-KRT 1-BGP RT Background
                AS path: I
                Communities: target:1:200 rte-type:0.0.0.0:1:0
                Import Accepted
                VPN Label: 299824 TTL Action: vrf-ttl-propagate
                Localpref: 100
                Router ID: 10.255.179.13
                Primary Routing Table bgp.l3vpn.0
```

show route table VPN-AB.inet.0

```
user@host> show route table VPN-AB.inet.0
      VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.39.1.0/30       *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
      10.39.1.4/30       *[Direct/0] 00:08:42
                          > via so-5/1/0.0
      10.39.1.6/32       *[Local/0] 00:08:46
                            Local
      10.255.71.16/32    *[Static/5] 00:07:24
                          > via so-2/0/0.0
```


Copyright © 2017, Juniper Networks, Inc.

```
10.255.71.17/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                     AS path: I
                    > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.18/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                     AS path: I
                    > via so-2/1/0.0, Push 100021, Push 100011(top)
10.255.245.245/32  *[BGP/170] 00:08:35, localpref 100
                     AS path: 2 I
                    > to 10.39.1.5 via so-5/1/0.0
10.255.245.246/32  *[OSPF/10] 00:07:24, metric 1
                    > via so-7/3/1.0
```

show route table VPN_blue.mvpn-inet6.0

```
user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65536:10.255.2.202/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                     > via so-0/1/3.0
1:10.255.2.203:65536:10.255.2.203/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
                      AS path: I
                     > via so-0/1/0.0
1:10.255.2.204:65536:10.255.2.204/432
                   *[MVPN/70] 00:57:23, metric2 1
                      Indirect
5:10.255.2.202:65536:128:::192.168.90.2:128:ffff::1/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                     > via so-0/1/3.0
6:10.255.2.203:65536:64500:128:::10.12.53.12:128:ffff::1/432
                   *[PIM/105] 00:02:37
                      Multicast (IPv6)
7:10.255.2.202:65536:64500:128:::192.168.90.2:128:ffff::1/432
                   *[MVPN/70] 00:02:37, metric2 1
                      Indirect
```

show route table vrf1.mvpn.0 extensive

```
user@host> show route table vrf1.mvpn.0 extensive
1:10.255.50.77:1:10.255.50.77/240 (1 entry, 1 announced)
        *MVPN   Preference: 70
                PMSI: Flags 0x0: Label 0: RSVP-TE:
Session_13[10.255.50.77:0:25624:10.255.50.77]
                Next hop type: Indirect
                Address: 0xbb2c944
                Next-hop reference count: 360
                Protocol next hop: 10.255.50.77
                Indirect next hop: 0x0 - INH Session ID: 0x0
                State: <Active Int Ext>
                Age: 53:03      Metric2: 1
                Validation State: unverified
                Task: mvpn global task
                Announcement bits (3): 0-PIM.vrf1 1-mvpn global task 2-rt-export

                AS path: I
```

Copyright © 2017, Juniper Networks, Inc.

show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 2
                State: <Active Ext>
                Local AS: 64502 Peer AS: 64500
                Age: 4
                Task: BGP_64500.10.12.99.5+3792
                Announcement bits (1): 0-Flow
                AS path: 64500 I
                Communities: traffic-rate:0:0
                Validation state: Accept, Originator: 10.12.99.5
                Via: 10.12.44.0/24, Active
                Localpref: 100
                Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
        *Flow   Preference: 5
                Next-hop reference count: 2
                State: <Active>
                Local AS: 64502
                Age: 6:30
                Task: RT Flow
                Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
                AS path: I
                Communities: 1:1

user@host> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.2:100:10.1.1.2/96 AD
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
10.1.1.4:100:10.1.1.4/96 AD
                   *[BGP/170] 1d 03:11:02, localpref 100, from 10.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
10.1.1.2:100:1:0/96 MH
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
10.1.1.4:100:1:0/96 MH
                   *[BGP/170] 1d 03:11:02, localpref 100, from 10.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
10.1.1.4:NoCtrlWord:5:100:100:10.1.1.2:10.1.1.4/176
                   *[VPLS/7] 1d 03:11:02, metric2 1
                    >    via ge-1/2/1.5
10.1.1.4:NoCtrlWord:5:100:100:10.1.1.4:10.1.1.2/176
                   *[LDP/9] 1d 03:11:02
                         Discard

user@host> show route table red extensive
red.inet.0: 364481 destinations, 714087 routes (364480 active, 48448 holddown, 1
 hidden)
10.0.0.0/32 (3 entries, 1 announced)
        State: <OnList CalcForwarding>
```


Copyright © 2017, Juniper Networks, Inc.

```
TSI:
KRT in-kernel 10.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

    Nexthop: Self
    AS path: [2] I
    Communities: target:2:1
Path 10.0.0.0 from 10.3.0.0 Vector len 4.  Val: 1
        @BGP    Preference: 170/-1
                Route Distinguisher: 2:1
                Next hop type: Indirect
                Address: 0x258059e4
                Next-hop reference count: 2
                Source: 2.2.0.0
                Next hop type: Router
                Next hop: 10.1.1.1 via ge-1/1/9.0, selected
                Label operation: Push 707633
                Label TTL action: prop-ttl
                Session Id: 0x17d8
                Protocol next hop: 10.2.0.0
                Push 16
                Composite next hop: 0x25805988 - INH Session ID: 0x193c
                Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
                State: <Secondary Active Int Ext ProtectionPath ProtectionCand>
                Local AS:     2 Peer AS:     2
                Age: 23         Metric2: 35
                Validation State: unverified
                Task: BGP_172.16.2.0.0+34549
                AS path: I
                Communities: target:2:1
                Import Accepted
                VPN Label: 16
                Localpref: 0
                Router ID: 10.2.0.0
                Primary Routing Table bgp.l3vpn.0
                Composite next hops: 1
                        Protocol next hop: 10.2.0.0 Metric: 35
                        Push 16
                        Composite next hop: 0x25805988 - INH Session ID: 0x193c
                        Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.1.1.1 via ge-1/1/9.0
                                Session Id: 0x17d8
                        2.2.0.0/32 Originating RIB: inet.3
                          Metric: 35                     Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 10.1.1.1 via ge-1/1/9.0
         BGP    Preference: 170/-1
                Route Distinguisher: 2:1
                Next hop type: Indirect
                Address: 0x9347028
                Next-hop reference count: 3
                Source: 10.3.0.0
                Next hop type: Router, Next hop index: 702
                Next hop: 10.1.4.2 via ge-1/0/0.0, selected
                Label operation: Push 634278
                Label TTL action: prop-ttl
                Session Id: 0x17d9
                Protocol next hop: 10.3.0.0
                Push 16
                Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da
                    State: <Secondary NotBest Int Ext ProtectionPath ProtectionCand>

                    Inactive reason: Not Best in its group - IGP metric
                    Local AS:     2 Peer AS:     2
                    Age: 3:34       Metric2: 70
                    Validation State: unverified
                    Task: BGP_172.16.3.0.0+32805
                    Announcement bits (2): 0-KRT 1-BGP_RT_Background
                    AS path: I
                    Communities: target:2:1
                    Import Accepted
                    VPN Label: 16
                    Localpref: 0
                    Router ID: 10.3.0.0
                    Primary Routing Table bgp.l3vpn.0
                    Composite next hops: 1
                            Protocol next hop: 10.3.0.0 Metric: 70
                            Push 16
                            Composite next hop: 0x93463a0 1048575 INH Session ID:
0x17da
                            Indirect next hop: 0x91e8800 1048574 INH Session ID:
0x17da
                            Indirect path forwarding next hops: 1
                                    Next hop type: Router
                                    Next hop: 10.1.4.2 via ge-1/0/0.0
                                    Session Id: 0x17d9
                            10.3.0.0/32 Originating RIB: inet.3
                              Metric: 70                      Node path count: 1
                              Forwarding nexthops: 1
                                    Nexthop: 10.1.4.2 via ge-1/0/0.0
                #Multipath Preference: 255
                        Next hop type: Indirect
                        Address: 0x24afca30
                        Next-hop reference count: 1
                        Next hop type: Router
                        Next hop: 10.1.1.1 via ge-1/1/9.0, selected
                        Label operation: Push 707633
                        Label TTL action: prop-ttl
                        Session Id: 0x17d8
                        Next hop type: Router, Next hop index: 702
                        Next hop: 10.1.4.2 via ge-1/0/0.0
                        Label operation: Push 634278
                        Label TTL action: prop-ttl
                        Session Id: 0x17d9
                        Protocol next hop: 10.2.0.0
                        Push 16
                        Composite next hop: 0x25805988 - INH Session ID: 0x193c
                       Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

                        Protocol next hop: 10.3.0.0
                        Push 16
                        Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
                      Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da Weight
0x4000
                        State: <ForwardingOnly Int Ext>
                        Inactive reason: Forwarding use only
                        Age: 23         Metric2: 35
                        Validation State: unverified
                        Task: RT
```


Copyright © 2017, Juniper Networks, Inc.

```
                    AS path: I
                    Communities: target:2:1
```

show route table bgp.evpn.0 extensive |no-more (EVPN)

```
                show route table bgp.evpn.0 extensive | no-more
                bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
                2:1000:10::100::00:aa:aa:aa:aa:aa/304 (1 entry, 0 announced)
                        *BGP    Preference: 170/-101
                                Route Distinguisher: 1000:10
                                Next hop type: Indirect
                                Address: 0x9420fd0
                                Next-hop reference count: 12
                                Source: 10.2.3.4
                                Protocol next hop: 10.2.3.4
                                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                         State: Local AS: 17  Peer AS:17  Age:21:12  Metric2:1  Validation State:
                 unverified
                                Task: BGP_17.1.2.3.4+50756
                                AS path: I
                                Communities: target:1111:8388708 encapsulation0:0:0:0:3
                                Import Accepted
                                Route Label: 100
                                ESI: 00:00:00:00:00:00:00:00:00:00
                                Localpref: 100
                                Router ID: 10.2.3.4
                                Secondary Tables: default-switch.evpn.0
                                Indirect next hops: 1
                                        Protocol next hop: 10.2.3.4 Metric: 1
                                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                                        Indirect path forwarding next hops: 1
                                                Next hop type: Router
                                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                                Session Id: 0x2
                                        1.2.3.4/32 Originating RIB: inet.0
                                          Metric: 1                       Node path count: 1
                                          Forwarding nexthops: 2
                                                Nexthop: 10.92.78.102 via em0.0

                2:1000:10::200::00:bb:bb:bb:bb:bb/304 (1 entry, 0 announced)
                        *BGP    Preference: 170/-101
                                Route Distinguisher: 1000:10
                                Next hop type: Indirect
                                Address: 0x9420fd0
                                Next-hop reference count: 12
                                Source: 10.2.3.4
                                Protocol next hop: 10.2.3.4
                                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                          State: Local AS:17  Peer AS:17  Age:19:43  Metric2:1  Validation
                State:unverified
                                Task: BGP_17.1.2.3.4+50756
                                AS path: I
                                Communities: target:2222:22 encapsulation0:0:0:0:3
                                Import Accepted
                                Route Label: 200
                                ESI: 00:00:00:00:00:00:00:00:00:00
                                Localpref: 100
                                Router ID: 10.2.3.4
                                Secondary Tables: default-switch.evpn.0
                                Indirect next hops: 1
                                        Protocol next hop: 10.2.3.4 Metric: 1
```

```
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

2:1000:10::300::00:cc:cc:cc:cc:cc/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
         State: Local AS:17  Peer AS:17  Age:17:21  Metric2:1  Validation State:
 unverified  Task: BGP 17,1,2,3,3,4+50756
          AS path: I
                Communities: target:3333:33 encapsulation0:0:0:0:3
                Import Accepted
                Route Label: 300
                ESI: 00:00:00:00:00:00:00:00:00:00
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::100::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                PMSI: Flags 0x0: Label 100: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
         State: Local AS:17  Peer AS:17 Age:37:01  Metric2:1  Validation State:
 unverified  Task: BGP 17.1.2.3.4+50756
                AS path: I
                Communities: target:1111:8388708 encapsulation0:0:0:0:3
                Import Accepted
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
```


Copyright © 2017, Juniper Networks, Inc.

```
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::200::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                PMSI: Flags 0x0: Label 200: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS: 17 Peer AS: 17   Age:35:22  Metric2:1  Validation
State:unverified  Task: BGP 17.1.2.3.4+50756
           AS path:I  Communities: target:2222:22 encapsulation):0:0:0:0:3

Import Accepted
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::300::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                PMSI: Flags 0x0: Label 300: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
         State: Local AS: 17 Peer AS:  17  Age 35:22 Metric2:1  Validation State:
 unverified   Task: BGP 17.1.2.3.4+5075
          6  AS path: I   Communities: target:3333:33  encapsulation0:0:0:0:3
Import Accepted   Localpref:100
          Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
```

```
        Next hop: 10.10.10.1 via xe-0/0/1.0
        Session Id: 0x2
10.2.3.4/32 Originating RIB: inet.0
  Metric: 1                       Node path count: 1
  Forwarding nexthops: 2
        Nexthop: 10.92.78.102 via em0.0
```


Copyright © 2017, Juniper Networks, Inc.

show route terse

List of Syntax
Syntax on page 383
Syntax (EX Series Switches) on page 383

Syntax
show route terse
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route terse

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display a high-level summary of the routes in the routing table.

NOTE: For BGP routes, the **show route terse** command displays the local preference attribute and MED instead of the metric1 and metric2 values. This is mostly due to historical reasons.

To display the metric1 and metric2 value of a BGP route, use the show route extensive command.

Options
none—Display a high-level summary of the routes in the routing table.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route terse on page 385

Output Fields
Table 36 on page 383 describes the output fields for the **show route terse** command. Output fields are listed in the approximate order in which they appear.

Table 36: show route terse Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number destinations	Number of destinations for which there are routes in the routing table.
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)

Copyright © 2017, Juniper Networks, Inc.

Table 36: show route terse Output Fields *(continued)*

Field Name	Field Description
route key	Key for the state of the route: • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • ***—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route.
A	Active route. An asterisk (*) indicates this is the active route.
V	Validation status of the route: • **?**—Not evaluated. Indicates that the route was not learned through BGP. • **I**—Invalid. Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **N**—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **V**—Valid. Indicates that the prefix and autonomous system pair are found in the database.
Destination	Destination of the route.
P	Protocol through which the route was learned: • **A**—Aggregate • **B**—BGP • **C**—CCC • **D**—Direct • **G**—GMPLS • **I**—IS-IS • **L**—L2CKT, L2VPN, LDP, Local • **K**—Kernel • **M**—MPLS, MSDP • **O**—OSPF • **P**—PIM • **R**—RIP, RIPng • **S**—Static • **T**—Tunnel
Prf	Preference value of the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
Metric 1	First metric value in the route. For routes learned from BGP, this is the MED metric.
Metric 2	Second metric value in the route. For routes learned from BGP, this is the IGP metric.

Copyright © 2017, Juniper Networks, Inc.

Table 36: show route terse Output Fields *(continued)*

Field Name	Field Description
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated.

Sample Output

show route terse

```
user@host> show route terse
inet.0: 10 destinations, 12 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 172.16.1.1/32        0  10          1            >10.0.0.2
  ?                     B 170        100                          I
  unverified                                      >10.0.0.2
* ? 172.16.1.1/32        D   0                       >lo0.2
* V 2.2.0.2/32          B 170        110                          200 I
  valid                                           >10.0.0.2
* ? 10.0.0.0/30         D   0                       >lt-1/2/0.1
  ?                     B 170        100                          I
  unverified                                      >10.0.0.2
* ? 10.0.0.1/32         L   0                        Local
* ? 10.0.0.4/30         B 170        100                          I
  unverified                                      >10.0.0.2
* ? 10.0.0.8/30         B 170        100                          I
  unverified                                      >10.0.0.2
* I 172.16.1.1/32       B 170         90                          200 I
  invalid                                         >10.0.0.2
* N 192.168.2.3/32      B 170        100                          200 I
  unknown                                         >10.0.0.2
* ? 172.16.233.5/32       0  10          1             MultiRecv
```

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.